

AD-A124 071


GENERALIZED MONITORING FACILITY USERS MANUAL(U) COMMAND
AND CONTROL TECHNICAL CENTER WASHINGTON DC
B WALLACK ET AL. 01 DEC 82 CCTC-CSM-UM-246-82-CHG-3


1/2

UNCLASSIFIED

F/G 9/2

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A



DEFENSE COMMUNICATIONS AGENCY
COMMAND AND CONTROL TECHNICAL CENTER
WASHINGTON, D. C. 20301

(12)

IN REPLY
REFER TO: C751

1 December 1982

TO: RECIPIENTS

SUBJECT: Change 3 to Computer System Manual CSM UM 246-82, Generalized
Monitoring Facility

ADA 124071

1. This is Change 3 to Computer System Manual CSM UM 246-82, - *AD A116898*
Generalized Monitoring Facility, dated 1 May 1982. Remove obsolete
pages and destroy them in accordance with applicable security
regulations and insert new pages as indicated below:

Remove Pages

viii
ix
xviii
xx
xxii
xxiii
xxiv
1-1 and 1-2
2-5 and 2-6
2-7 and 2-8
2-9 and 2-10
2-11 and 2-12
2-12.1 and 2-12.2
5-1 and 5-2
5-3 and 5-4
5-5 and 5-6
5-7 and 5-8
5-9 and 5-10
5-11 and 5-12
5-15 and 5-16
5-17 and 5-18
5-23 and 5-24
5-27 and 5-27.1
5-27.18 and 5-27.19
5-31 and 5-32
5-55 and 5-56
6-3 and 6-4
6-5 and 6-6
6-7 and 6-8
6-15 and 6-16
6-19 and 6-19.1
6-23 and 6-24

Insert New Pages

viii
ix
xviii
xviii.1
xviii.2
xx
xxii
xxiii
xxiii.1
xxiv
1-1 and 1-2
2-5 and 2-6
2-7 and 2-8
2-9 and 2-10
2-11 and 2-12
2-12.1 and 2-12.2
5-1 and 5-2
5-3 and 5-4
5-5 and 5-6
5-7 and 5-8
5-9 and 5-10
5-11 and 5-11.1
5-11.2 and 5-12
5-15 and 5-16
5-17 and 5-18
5-23 and 5-24
5-27 and 5-27.1
5-27.18 and 5-27.19
5-31 and 5-32
5-55 and 5-56
6-3 and 6-4
6-5 and 6-6

DTIC
SELECTED
FEB 03 1983
E

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

DTIC FILE COPY

83 02 02 005

Remove Pages

6-25 and 6-26
6-37 and 6-38
6-39 and 6-40
6-53 and 6-53.1
6-53.2 and 6-54
6-55 and 6-56
6-59 and 6-60
9-23 and 9-24
9-25 and 9-26
14-13 and 14-14
14-17 and 14-18
14-21 and 14-22
14-23 and 14-24
15-1 and 15-2
16-1 and 16-2


Insert New Pages

6-6.1 and 6-6.2
6-7 and 6-8
6-15 and 6-15.1
6-15.2 and 6-16
6-19 and 6-19.1
6-23 and 6-24
6-25 and 6-26
6-37 and 6-38
6-39 and 6-40
6-53 and 6-53.1
6-53.2 and 6-54
6-55 and 6-56
6-59 and 6-60
9-23 and 9-24
9-25 and 9-26
14-13 and 14-14
14-17 and 14-18
14-18.1 and 14-18.2
14-21 and 14-22
14-22.1 and 14-22.2
14-23 and 14-24
15-1 through 15-64
(new section)
16-1 and 16-2
17-1 and 17-2

2. The effective date of these change pages is 1 December 1982.
3. When this change has been made, post an entry in the Record of Changes and file this letter before the title page.

FOR THE DIRECTOR:

87 Enclosures


FOI ALAN GROSS
Colonel, USAF
Deputy Director
for Computer Services

Section		Page
6.3.3.39	Report 50 - The Original Allocation Time for User Memory in 1/10 Second	6-41
6.3.3.40	Report 51 - The Time Corrected Percent of Assigned Memory Used	6-41
6.3.4	Activity Resource Usage Report	6-41
6.3.5	SNUMB-IDENT Report	6-43
6.3.6	Memory Map Report	6-43
6.3.7	Demand List Report	6-48
6.3.8	Activity Abort Report	6-48
6.3.9	Jobs Out of Core Report	6-50
6.3.10	Excessive Resource Use Report	6-50
6.3.11	Allocation Status Report	6-53
6.3.12	Plots Reports	6-53
6.3.12.1	Plot 1 - Available Memory vs. Outstanding Demand in Core Allocator Queue vs. Outstanding Demand in Core Allocator + Peripheral Allocator Queue	6-56
6.3.12.2	Plot 2 - Memory Shortfall in Core Allocator vs. Memory Shortfall in Core Allocator + Peripheral Allocator	6-56
6.3.12.3	Plot 3 - Number of Activities in Core Queue vs. Number of Activities in Peripheral Allocator Queue	6-56
6.3.12.4	Plot 4 - Average Size of TSS, FTS and NCP	6-56
6.3.13	Memory Statistics Report	6-58
6.3.14	Special Job Memory Reports	6-58
6.3.15	Distribution of Urgency Over Time Report	6-58
6.3.16	Zero Urgency Job Report	6-58.3
6.4	Error Messages	6-58.4
6.5	Multireel Processing	6-62
6.6	TAPE Error Aborts	6-63
7.	MASS STORE DATA REDUCTION PROGRAM (MSDRP)	7-1
7.1	Introduction	7-1
7.2	Data Collection Methodology	7-3
7.3	Analytical Methodology	7-3
7.4	Data Reduction Methodology	7-4
7.5	MSMDRP Output	7-5
7.5.1	System Configuration and Channel Usage Report (File 42)	7-5
7.5.2	System Summary Report (File 42)	7-9
7.5.3	System Traces Captured by Monitor Report (File 42)	7-10
7.5.4	Channel Status Changes Report (File 29)	7-10
7.5.5	Physical Device, Device ID Correlation Table (File 42)	7-10
7.5.6	Device Space Utilization Report (File 42)	7-10
7.5.7	Device Seek Movement Report (File 42)	7-15
7.5.8	Head Movement Efficiency Report (File 42)	7-17

Section		Page
14.4.2.2.6	Operations Practices	14-7
14.4.2.2.7	Batch Job Scheduling	14-7
14.4.2.2.8	Site and Computer Facility Organization . . .	14-7
14.4.2.3	Understand Installation Service Objectives and Priorities	14-7
14.4.2.3.1	Installation Service Objectives	14-7
14.4.2.3.1.1	Production Objective	14-8
14.4.2.3.1.2	Availability Objective	14-8
14.4.2.3.1.3	Mixed Objectives	14-8
14.4.2.3.2	Installation Priorities	14-9
14.4.2.3.2.1	Service Priorities	14-9
14.4.2.3.2.2	Evaluation/Tuning Solution Priorities . . .	14-9
14.4.2.4	Specify Current Evaluation	14-9
14.4.2.4.1	Objective	14-9
14.4.2.4.2	Objective Decision	14-9
14.4.2.4.2.1	Attainable Objectives	14-9
14.4.2.4.2.2	Realistic Objectives	14-9
14.4.2.4.2.3	Cost-Effective Objectives	14-10
14.4.2.4.3	Determine if Worth Continuing	14-10
14.4.2.4.4	Begin Problem Analysis	14-10
14.4.3	Problem Analysis Phase	14-10
14.4.3.1	Run Appropriate Analysis Tool	14-10
14.4.3.2	Evaluate System Output	14-10
14.4.3.3	Follow Tuning Procedures	14-10
14.4.3.4	Evaluate Need to Continue the Tuning	14-10
14.5	Composition of a Performance Evaluation Team	14-11
14.6	System Evaluation	14-11
14.6.1	Introduction	14-11
14.6.2	Selecting a Representative Value From GMC Histogram Reports	14-11
14.6.2.1	Symmetric Distribution Closely Clustered Around a Single Point	14-11
14.6.2.2	Skewed Distribution	14-13
14.6.2.3	Distribution With Outliers	14-13
14.6.3	Memory Evaluation	14-13
14.6.3.1	Obtaining the Data	14-15
14.6.3.2	Evaluating the Data	14-15
14.6.4	CPU Evaluation	14-18.1
14.6.4.1	Data Recording	14-18.1
14.6.4.2	Evaluating the Data	14-19
14.6.5	I/O Evaluation	14-22
14.6.5.1	Data Recording	14-22
14.6.5.2	Evaluating the Data	14-22
14.6.5.3	Mass Storage Operation	14-25

Section		Page
15.	TIMESHARING EMULATION AND RESPONSE SYSTEM	
	(TEARS)	15-1
15.1	Introduction	15-1
15.2	Data Collection Methodology	15-1
15.3	Analytical Methodology	15-2
15.4	Data Reduction Methodology	15-3
15.4.1	Response Pass Reduction Methodology	15-4
15.4.2	Emulation Reduction Methodology	15-5
15.5	Output	15-7
15.5.1	TEARS Response Pass Reports	15-38
15.5.1.1	TSS Reduction Event Log	15-38
15.5.1.2	Response Time Logplots	15-41
15.5.1.3	Subdispatch Time Logplots	15-41
15.5.1.4	CPU Times Allocated to TSI-Exec and TSI-Subdispatch Report	15-42
15.5.1.5	TSS Subtraces Encountered, and Derails Executed by Users	15-43
15.5.2	TEARS Emulation Reports	15-43
15.5.2.1	Exception Message Report	15-44
15.5.2.2	Users' Map	15-44
15.5.2.3	TSS Delay - Users Delay Report	15-45
15.5.2.4	TSS Core Map	15-47
15.5.2.5	Error Message Report	15-48
15.5.2.6	Intertrace Gap - Duration	15-49
15.5.2.7	Memory Activity Report	15-49
15.5.2.8	Eligibles for Subdispatch - Duration	15-50
15.5.2.9	In Subdispatch - Duration	15-50
15.5.2.10	Eligibles for and In Subdispatch - Duration	15-50
15.5.2.11	User Swaps - Swap Rate	15-50
15.5.2.12	User Swaps - Swap Amount	15-50
15.5.2.13	User Swaps - Duration	15-51
15.5.2.14	User I/Os - Duration	15-51
15.5.2.15	In System Master Catalog Wait - Duration	15-51
15.5.2.16	Intertrace Time (Adjusted for TSS GEWAKES)	15-51
15.5.3	Formatted Dump Reports	15-51
15.5.3.1	Formatted Dump	15-51
15.5.3.2	Intertrace Time	15-52
15.6	Reduction Run-Time Options	15-52
15.6.1	Parameter List	15-53
15.6.2	Command Syntax	15-56
15.6.3	Examples	15-61
15.6.4	Card Input Errors	15-61
15.7	JCL for Reduction Execution	15-62
15.8	Multireel Processing	15-62
15.9	Tape Error Aborts	15-64

Section

Page

DISTRIBUTION	16-1
DD Form 1473	17-1

Figure		Page
6-3	Standard Plot	6-13
6-4	PLOT Action Code Format	6-14
6-5	TIME Action Code Format	6-16
6-6	ALLOFF/ALLON Action Code Format	6-17
6-7	System Bottleneck Chart	6-20
6-8	JCL to RUN MUDRP	6-22
6-9	MUM Title Page Report - Idle Monitor Active	6-25
6-10	MUM Title Page Report - Idle Monitor Off	6-27
6-11	System Program Load	6-30
6-12	Standard Histogram Report	6-31
6-13	Out-of-Range Histogram	6-33
6-14	Report 25 -- THIS FIGURE DELETED	6-39
6-15	Activity Resource Usage Report	6-42
6-16	SNUMB IDENT Report	6-44
6-17	Memory Map Report	6-45
6-18	Demand List Report	6-49
6-19	Abort Report	6-51
6-20	Jobs Out of Core Report	6-52
6-21	Excessive Resource Usage Report	6-54
6-22	Allocation Status Report	6-55
6-23	Standard Plot	6-57
6-24	Memory Statistics Report	6-59
6-25	Special Job Memory Demand Report	6-60
6-26	Special Job Memory Size Report	6-61
6-27	Distribution of Urgency Over Time Report	6-61.1
6-28	Zero Urgency Job Report	6-61.2
7-1	System Configuration and Channel Usage Report	7-6
7-2	MSM System Summary Report	7-8
7-3	System Traces Captured by Monitor Report	7-11
7-4	Channel Status Changes Report	7-12
7-5	Physical Device, Device ID Correlation Table	7-13
7-6	Device Space Utilization Report	7-14
7-7	Device Seek Movement Report	7-16
7-8	Head Movement Efficiency Report	7-18
7-9	System File Use Summary Report	7-20
7-10	Individual Module Activity Report	7-22
7-11	SSA Module Usage Report by Job	7-24
7-12	File Code Summary Report	7-25
7-13	Cat/File String Report	7-27
7-14	Connect Summary Report by USERID/SNUMB	7-30
7-15	Activity Summary Report	7-31
7-16	Device Area File Code Reference Report	7-33
7-17	Device File Use Summary Report	7-34
7-18	Chronological Device Utilization Report	7-35
7-19	FMS Cache Report	7-37
7-20	Proportionate Device Utilization Report	7-38

Figure		Page
10-1	Default Values For Plots	10-5
10-2	Average Response Time for All Uses by DATANET (Report RESPD)	10-10
10-3	HSLA/Subchannel Transmission Report (Report BUFF)	10-13
10-4	Average Response Time for Designated Lines (Report RESPL)	10-14
10-5	HSLA Subchannels Being Monitored (Report HSLA)	10-15
10-6	Card Images From I* FILE of the GMF Monitor (Report CARD)	10-17
10-7	List of Active Line IDs	10-18
10-8	HSLA Threshold Report (Report Thresh)	10-19
10-9	Number of Transactions Send to Host (Plot 1)	10-23
10-10	Number of 36 Bit Words Received From Host (Plot 4)	10-25
10-11	Percent Idle Time (Plot 5)	10-26
10-12	Percent Buffer Requests Denied (Plot 6)	10-27
10-13	Histogram of DATANET Response	10-29
10-14	Histogram Special Line ID Response	10-30
10-15	Plot Alteration (Action Code PLOT)	10-34
10-16	Set Timespan(s) For Processing (Action Code TIME) . . .	10-36
10-17	Action Codes For Turning Reports Off	10-37
10-18	Action Codes For Turning Reports On	10-38
10-19	Turn Off Plots (Action Code NOPLOT)	10-41
10-20	Threshold Values (Action Code THRESH)	10-43
10-21	Turn Off Reports For Selected DATANETs (Action Code NORPT)	10-44
10-22	Debug Option (Action Code DEBUG)	10-46
10-23	Alter Histogram Parameters (Action Code CHANGE)	10-47
10-24	Unmatched/Duplicate Error Report (Action Code DUPLIC)	10-48
10-25	Reject Report (Action Code REJECT)	10-49
11-1	Sample Execution Report	11-5
11-2	CPU Title Page	11-8
11-3	Number of System Activities in the CPU Queue	11-9
11-4	Number of User Activities in the CPU Queue	11-10
11-5	CPU Burst Length Distribution for User Activities . . .	11-11
11-6	CPU Burst Length Distribution for All Activities . . .	11-13
11-7	Number of System Activities in Core With Outstanding I/O	11-14
11-8	Number of User Activities in Core With Outstanding I/O	11-15
11-9	Number of System Activities in Core With CPU-I/O Overlap	11-16
11-10	Number of User Activities in Core With CPU-I/O Overlap	11-17
11-11	Number of Inactive Activities in Core	11-18

Figure

		Page
11-12	CPU Time Report	11-20
11-13	CPU Time Report With Special Snumbs	11-22
11-14	CPU Plot of Percent Idle	11-23
11-15	WIN Report	11-24
11-16	Tape Reduction Title Page	11-25
11-17	Number of Tape Drives in Use (Time Corrected)	11-27
11-18	Tape Activity Report	11-28
11-19	Tape Status Report	11-29
11-20	Sample Deck Setup	11-40
12-1	TPS GMF Compile JCL	12-4
12-2	Control Cards for Running Only the TPS GMF	12-5
12-3	JCL for TPE Data Reduction	12-7
12-4	TPE Summary Report	12-9
12-5	NONEW Option Report	12-10
12-6	Profile ID Report	12-11
12-7	Buffer Space Report	12-12
12-8	Table Reaching Threshold Size Report	12-14
12-9	Master Terminal Report	12-15
12-10	Impasse Report	12-17
12-11	Output Transmission Report	12-18
12-12	TPAP ABORT Report	12-19
12-13	UNDOT Entry Report	12-20
12-14	Reissuing Intercom Message Report	12-21
12-15	Receiving Intercom Message Report	12-22
12-16	TPAP Modify Report	12-23
12-17	TPAP Status Report	12-24
12-18	TPAP Run Times Report	12-25
12-19	JCL For TPE Formatted Dump	12-28
12-20	TPE Formatted Dump	12-29
12-21	Transaction Processing System Alter Cards	12-31
13-1	New Initialization Section	13-2
13-2	Monitor Subroutine	13-3
13-3	Creation of User Trace	13-5
14-1	Flowchart of Evaluation/Tuning Process	14-4
14-2	Sample Symmetric Distribution	14-12
14-3	Sample Skewed Distribution	14-14
14-4	Memory Statistics	14-16
14-5	CPU Program Characteristics	14-20
14-6	CPU Processor Availability	14-21
14-7	Temporary Storage Test Form	14-24
15-1	Execution Summary Report	15-8
15-2	Timesharing Reduction Event Log Report	15-9
15-3	Response Time For All Users Report	15-11
15-4	CPU Times Allocated to Exec and Subdispatch Report	15-12
15-5	Response Time For Users Not Requesting More Core Plot	15-13

Figure		Page
15-6	Response Time For Users With Core Requests	
	During Line Idle Plot	15-14
15-7	Total Time In Subdispatch Queue Plot	15-15
15-8	Total Time In Subdispatch Queue Waiting	
	Service Plot	15-16
15-9	Processor Time In Subdispatch Plot	15-17
15-10	TSS Subtraces Encountered Histogram	15-18
15-11	Derails Executed By User Histogram	15-20
15-12	Execution Message Report	15-22
15-13	Users Map	15-23
15-14	TSS Delay and Users Delay Report	15-24
15-15	TSS Core Map	15-25
15-16	Error Message Report	15-26
15-17	Intertrace Gap Duration Plot	15-27
15-18	Memory Activity Plot	15-28
15-19	Eligibles For Subdispatch Plot	15-29
15-20	In Subdispatch Plot	15-30
15-21	Eligibles For and In Subdispatch Plot	15-31
15-22	User Swap Rate Plot	15-32
15-23	User Swap Amount Plot	15-33
15-24	User Swap Duration Plot	15-34
15-25	User I/OS Duration Plot	15-35
15-26	In System Master Catalog (SMC) Wait Plot	15-36
15-27	Intertrace Time (Adjusted For GEWAKES)	
	Histogram	15-37
15-28	Formatted Dump Report	15-39
15-29	Intertrace Time Histogram	15-40
15-30	TEARS JCL	15-63

TABLES

Number		Page
2-1	GMC Release Dependent Parameters	2-11
5-1	Required Trace Type for Each Monitor	5-5
5-2	Abort Codes	5-6
5-3	GMC Catalog and File Index	5-9
6-1	Default Reports	6-2
6-2	Default Values for Histograms	6-6
6-3	Default Values for Plots	6-7
6-4	Available Report Actions and Their (Default) Values . .	6-8
6-5	File Code for MUM Reports	6-24
7-1	MSM/MSMDRP Reports	7-2
9-1	Octal Codes and Device Names for the Mailbox and Statistical Summary Reports	9-9
9-2	Acceptable Device Type Names	9-34
10-1	DATANET-355 Reports	10-6
10-2	Response Time Error Messages	10-21
11-1	Central Processing Unit and Tape Reduction Reports . .	11-2
11-2	Histogram Default Parameters	11-30
11-3	Plot Default Parameters	11-31
12-1	Transaction Processing System Traces	12-2
15-1	Emulation Phase User States	15-6
15-2	Histogram Default Parameters	15-54
15-3	Plot Default Parameters	15-55

SECTION 1. GENERAL

1.1 Purpose of the Users Manual

The Users Manual describes each of the programs in the Generalized Monitoring Facility (GMF), discusses input options required to run each program, and provides sample outputs generated by each program.

The Generalized Monitoring Facility is delivered on a FILSYS SAVE tape. A description of the software is presented in section 2. Installation procedures for the GMF Monitor and guidance on how to run GMF can be found in section 5 of this manual.

1.2 Project References

The Generalized Monitoring Facility was originally developed for the Government by Honeywell Information Systems. Since delivery of the completed software in 1975, the Computer Performance Evaluation Branch (C751) of the Command and Control Technical Center has extensively modified and rewritten the GMF system. Numerous software errors have been corrected and many new features have been added.

1.3 Terms and Abbreviations

The following abbreviations will be used throughout the document.

CAM	-	Communications Analysis Monitor
CM	-	Channel Monitor
CPU	-	Central Processing Unit
CPUM	-	CPU Monitor
GCOS	-	Generalized Comprehensive Operating System
GMC	-	Generalized Monitoring Collector
GMF	-	Generalized Monitoring Facility
GRTS	-	General Remote Terminal Supervisor
GRTM	-	GRTS Monitor
IDLEM	-	Idle Monitor

- MSM - Mass Storage Monitor
- MUM - Memory Utilization Monitor
- RMC - Resource Monitor Collector
- RMDRx - Resource Monitor Data Reduction Program 1 through 3
- RMON - Resource Monitor
- TM - Tape Monitor
- TPEM - Transaction Processing Executive Monitor
- TSSM - Timesharing Subsystem Monitor
- WWMCCS - World Wide Military Command and Control System

1.4 Security and Privacy

There are no classified data collected in the GMF, but there is one exception. If the Communications Analysis Monitor (CAM) is being run with the Specific Terminal Option, classified data may be collected.

1.5 Manual Format

Section 2 of this Manual provides an overview of the GMF system. A brief description of each program is included. Sections 3 through 12 and 15 describe the programs of the GMF system in detail. In the presentation the user will find the appropriate information to successfully operate each program. The format of the sections includes a discussion of each program in the input, processing, and output phases. Section 13 of the document describes the procedures that a user should follow if it is desired to create a new GMF monitor and section 14 provides detailed information as to how GMF should be used in conducting a complete system-wide evaluation.

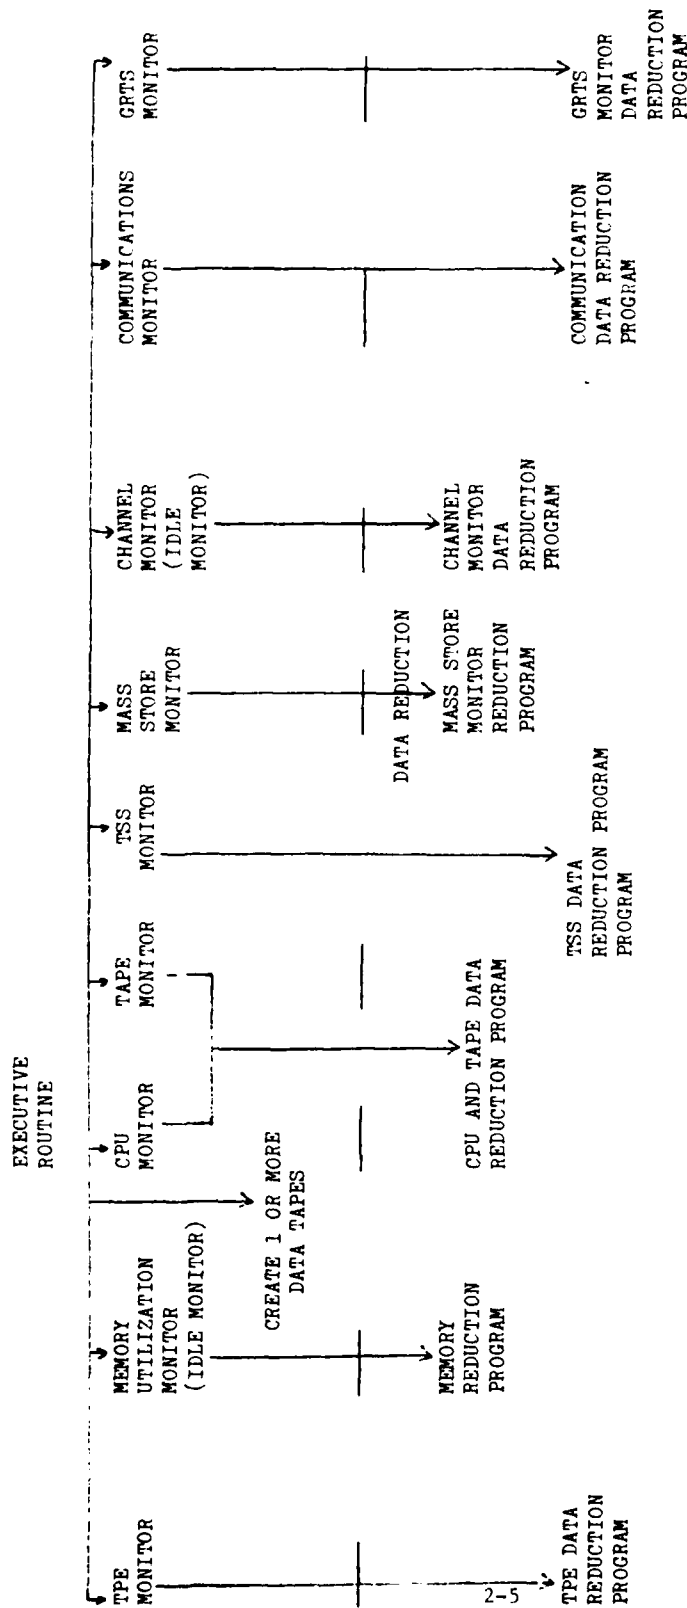


Figure 2-3. Programs in the GWC Subsystem

<u>Data Collector Programs</u>	<u>Subroutines</u>	<u>Traces Captured</u>
Memory Utilization Monitor	T10 T46	10,11,51 46
Idle Monitor	T21 TRCS	21 0,1,2,3,13,16,22, 37,65
Mass Store Monitor	T7	7,15,73*,76*,77*
Channel Monitor	T4,T7,T22	4,7,15,22
Tape Monitor	T50	50,51,52
CPU Monitor	T70	10,11,21,70*
Communications Analysis Monitor	T14	14*
GRTS Monitor	T62	62*
TPE Monitor	T200	0,1,2,4,5,6,13, 42,51,65,74*
TSS Monitor	T100	74*

* - Nonstandard traces generated by the particular monitor.

Figure 2-4. Subroutines and Traces in GMC Data
Collector Programs

2.4 System Organization

The GMF is composed of two data collectors, GMC and RMC, and associated data reduction programs. Sections 3 through 12 describe these programs. Figure 2-1 gives a system flowchart for the RMC. Figure 5-1 gives a system flowchart for the GMC.

2.5 Performance

The GMF monitors the performance of a system, aids in identifying the start of system performance problems, and aids in analyzing system performance problems. The RMC requires very little system resource usage and writes all its data to the system accounting file. The GMC is a much more detailed system with the associated higher overhead. The GMC is used mainly to determine the cause of system performance problems. The GMC requires 15 to 24 thousand words of memory and one tape drive while being run. Both systems require offline data reduction.

2.6 GMF Installation

2.6.1 Creation of GMF Files. The GMF software is contained on a single user save tape. The USERID on the tape is B29IDPX0. This USERID must be created with 3950 LLINKS of file space. A user restore can then be run. B29IDPX0 is subdivided into several catalogs described below:

- . GMFCOL - 530 LLINKS - This subcatalog contains all the data collection software for the GMC monitoring system. All files within this subcatalog are completely described in section 5.

- . SOURCE - 2450 LLINKS - This subcatalog contains Time Sharing source files for all data reduction programs contained within the GMC system. Figure 2-5 is a breakdown of the individual files within this subcatalog. Sections 6-12 describe each program in detail.

- . OBJECT - 1390 LLINKS - This subcatalog contains the object decks for all the data reduction programs contained within the GMC system. Figure 2-5 is a breakdown of the individual files within this subcatalog.

- . JCL - 16 LLINKS - This subcatalog contains all the JCL required to run all the data reduction programs contained within the GMC system. Figure 2-6 is a breakdown of the individual files within this subcatalog.

- . RMON - 345 LLINKS - This subcatalog contains all the software required to collect and reduce the data for the RMON Monitoring system. This subcatalog is further subdivided into JCL, SOURCE and OBJECT subcatalogs. The files within these subcatalogs are completely described in sections 3 and 4.

<u>FILE NAME</u>	<u>FUNCTION</u>	<u>SOURCE SIZE (LL)</u>	<u>OBJECT SIZE (LL)</u>
MUM	MEMORY UTILIZATION MONI- TOR DATA REDUCTION PRO- GRAM. REFERENCED IN SECTION 6.	375	220
MSM	MASS STORE MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 7.	320	190
CM	CHANNEL MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 8.	254	182
CAM	COMMUNICATION MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 9.	184	110
CPU-TAPE	CPU AND TAPE MONITOR DATA REDUCTION PROGRAMS. REFERENCED IN SECTION 11.	390	155
GRT	DATANET 355 MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 10.	280	154
TPETG	TRANSACTION PROCESSING DATA REDUCTION PROGRAM. REFERENCED IN SECTION 12.	64	71
TPEALT	AN ALTER FILE FOR ADDING TPE TRACE CODE INTO THE TPE SUBSYSTEM (NO OBJECT FILE). REFERENCED IN SECTION 12.	14	
TPEDUMP	A PROGRAM FOR OBTAINING A FORMATTED TRACE DUMP FROM A TPE/GMF DATA TAPE. REFERENCED IN SECTION 12.	38	26
TSS	TIMESHARING MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 15.	530	282

Figure 2-5. B29IDPXO/SOURCE and
B29IDPXO/OBJECT Catalog Structure

<u>FILE NAME</u>	<u>FUNCTION</u>	<u>SOURCE SIZE (LL)</u>
MUM	JCL TO OBTAIN ALL MEMORY UTILIZATION MONITOR REPORTS	2
MSM	JCL TO OBTAIN MASS STORE MONITOR REPORTS	2
CM	JCL TO OBTAIN CHANNEL MONITOR REPORTS	2
CAM	JCL TO OBTAIN COMMUNI- CATION MONITOR REPORTS	2
GRT	JCL TO OBTAIN DN-355 MONITOR REPORTS	2
CPU-TAPE	JCL TO OBTAIN CPU AND TAPE MONITOR REPORTS	2
TPETG	JCL TO OBTAIN ALL REPORTS FROM THE TPE DATA REDUCTION PROGRAM	2
TSS	JCL TO OBTAIN ALL REPORTS FROM THE TSS DATA REDUCTION PROGRAM	2

Figure 2-6. B29IDPX0/JCL Catalog Structure

2.6.2 GMF Release Dependent Parameters. In order for the GMC to operate properly, it is necessary for GMC to locate certain instructions and/or words within several system programs. The user should insure that these locations are correct for the particular GCOS release, under which he is operating. Table 2-1 is a list of these dependent parameters identifying their the use and providing the approximate program source line numbers where the particular parameters are used. The list is provided for each GMC program that must checked by the user.

In order to use the GMC data reduction programs on a WW6.4 system, there is a special data card required in certain programs. This option applies to all data reduction programs except CPU-TAPE and TPEIG. The TPEIG program is not designed for use under any release other than WW7.2. The CPU-TAPE program would require a one-line source change to be used under a WW6.4 release.

The GMC system is designed so that data collected on a WW6.4 system may be reduced under a WW6.4 system or a WW7.2 system. In addition, data collected under a WW7.2 system may likewise be reduced under a WW7.2 system, or a WW6.4 system. Whenever the data reduction programs for MUM, CM, CAM, or GRT are used on a WW6.4 system, a data card with an RN typed on it must be included in the input section of the JCL deck. It makes no difference under what release the data was collected. It is only a question of under what release the data is being reduced.

For the MSM data reduction program, there are two data cards required. The first data card always contains the letters RN. The second card is determined by the following table:

<u>Data Collected</u>	<u>Data Reduced</u>	<u>Data Value</u>
WW6.4	WW6.4	1
WW6.4	WW7.2	3
WW7.2	WW6.4	2
WW7.2	WW7.2	NO SPECIAL CARDS REQUIRED

For the CPU-TAPE data reduction program, it is required to delete source line #4320 and recompile when using the data reduction program under a WW6.4 release.

The RMC System is also designed to run under GCOS release WW6.4.2 (commercial release 2H), or WW7.2 (commercial release 4JS (any level)). See subsection 3.3.1 for details as to required modifications.

Table 2-1. GMC Release Dependent Parameters

<u>Program</u>	<u>LINE #</u>	<u>Variable</u>	<u>Explanation</u>
GMF.TOP	90	SYS64	Used to control conditional assembly of GMC set=1 for W6.4(2H) release set=0 for W7.2(4J) release
	10220	-	Code in this area searches for trace processing within the dispatcher. Trace code must be within 500 octal locations of the address specified by entry pt 15 decimal of the dispatcher. The code being searched for is a LDAQ;STAQ;TRA0,1
	10700	-	Code in this area is used to make a correction to accounting processing, if the correction has not already been made via patches. The code is searched for within 500 octal locations of .MIOS entry point. The code searched for is SBLA TRREG+7,\$;ARL 12; ADLA .CRTOD,7. The ARL is changed to an ARS.
CPU.PAT	10	-	Code in this area searches for an ASA .SALT,5 Instruction in the dispatcher.
	210		In W6.4 search octal locations 1500-1700.
	230		In W7.2 search octal locations 2340-2450. In addition in W7.2 we need to find a STQ .QTOD,4 instruction
	420		between octal locations 2400-2460. For both releases we need to find 8 words of patch space. In W6.4
	720		between octal locations 3540-3740. In W7.2 between octal locations
	740		4600-5000. If not found here then
	1460		search octal locations 4150-4300 in
	1480		W6.4 or octal locations 5400-5530 in W7.2 all offsets are from the entry point of .MDISP. In addition .MFALT is searched in W7.2 for an ARL 12 instruction
	1190		between octal locations 2500-2550 (offset from the entry point). This is for gate locked timing code which is supposed to be assembled into W7.2 code.

<u>Program</u>	<u>LINE #</u>	<u>Variable</u>	<u>Explanation</u>
CAM.INIT	350	-	Beginning at 1400 octal locations from the entry point of .MDNET and continuing for 5000 octal locations search for a ANA=0777777,DL (777777375207) instruction, followed by a CMPA (0000000115210) instruction. This searches for # special interrupts processed code (NSIP).
	510	-	Beginning at the CMPA location found above in .MDNET and continuing for 3000 octal locations search for a ANA=077 followed by a CMPA=077. When found, back up 30 octal words and look for an AOS instruction. This searches for the # of lines waiting mailbox code (ROLXCT).
CAM.PAT	10	-	Code in this area searches for a LDQ M.LID,3 instruction in DNWW, followed by a ANQ =0077777,DU instruction. Octal
	140		locations 5000-6000 are searched (offsets are from entry point). Ten words of patch space must also be obtained. This patch
	370		space must be between octal locations 11100 and 11200 (offsets are from entry point). If patch space is not found then 7 words of patch space are searched for within the dispatcher. This search is performed between octal locations
	720		3540-3740 in W6.4 and octal locations
	740		4600-5000 in W7.2 (offsets are from entry point). It should be noted that in commercial releases and WW7.2, DNWW is referred to as DNET.
MSM.PAT	10	-	Code in this area searches for an AOS .CRTDL instruction and an AOS .CRTBH instruction in the dispatcher. In W6.4
	390-680		and W7.2, these need to be within 300 octal locations of the label DBASE. If these instructions are not found, a search is made from octal locations 4600-5100 in W6.4 and octal locations 7164-7464 in W7.2. In addition, 8 words of patch space is needed. In W6.4 between octal
	850		3540-3740. In W7.2, between octal
	870		4600-5000. If the patch space is not found then search between octal locations
	1320		4150-4300 in W6.4, or octal 5400-5530 in

<u>Program</u>	<u>LINE #</u>	<u>Variable</u>	<u>Explanation</u>
	1340		W7.2. All offsets are from the entry point of .MDISP. In addition in W7.2 FMS CACHE logic is also analyzed. See label TSFIO in routine T7 for locations required within FSIO module.
GMF.MON	840	FMS1	Offset from entry point of .MFSIO which points to the word giving the absolute address of FMS catalog cache buffer. Used only in W7.2. Set to -13 decimal.
	850	FMS2	Offset from entry point of .MFSIO pointing to the work which gives the option selection for FMS catalog cache. Used only in W7.2. Set to -15 decimal.

THIS PAGE LEFT INTENTIONALLY BLANK

SECTION 5. THE GENERAL MONITOR COLLECTOR - DATA COLLECTION PROGRAM

5.1 Introduction

The General Monitor Collector (GMC) data collection program is a privileged slave program which processes GCOS system trace data, organizes the data in GMC records, and writes the collected GMC data records on its output magnetic tape. The general concept of operation for the GMC facility is shown in figure 5-1.

As a privileged slave program, the monitor data collector requires the permission of the system operator to run and must execute in master mode. The master mode capability allows the GMC to access all of the system main memory. The areas of interest are the system Communication Region (CR) and the individual job Slave Service Areas (SSAs).

The GMC is actually a series of independent data collection monitors which are controlled via the central Executive Routine (ER) (described in subsection 5.3.1), and which use a common buffering routine for writing collected data to a common tape. The current GMC is comprised of 10 different monitors, which can be executed independently, or in any combination, except that TPE and TSS monitors cannot be run simultaneously. The monitors are described in this section. Each of the monitors has a dedicated data reduction program that produces formatted reports. These data reduction programs are described in sections 6 through 12 and 15.

The GMC can obtain control from the system in one of three ways. First, the standard manner is for the GMC to obtain control via the normal system trace. Second and third, GMC may also gain control in two nonstandard manners. These are via internally created system sub-traces (referred to as IT traces) and direct transfer patches (referred to as DT traces).

In the first way, the standard mechanism used by the monitor for obtaining control from the operating system is the normal GCOS system trace. The system trace capability records a history of the occurrence of as many as 72 system events, 64 of which are presently defined. This recording takes place in a circular table in the Communication Region of memory and is accomplished by execution of a unique code set resident in the system Dispatcher Module (.MDISP). Execution of this code is common to all system trace events and provides the point at which the GMC obtains control. The initialization portion of the GMC locates the system trace code set and implants a transfer instruction to the GMC executive. Thus, whenever a system trace event occurs anywhere within the system, the GMC executive will obtain control.

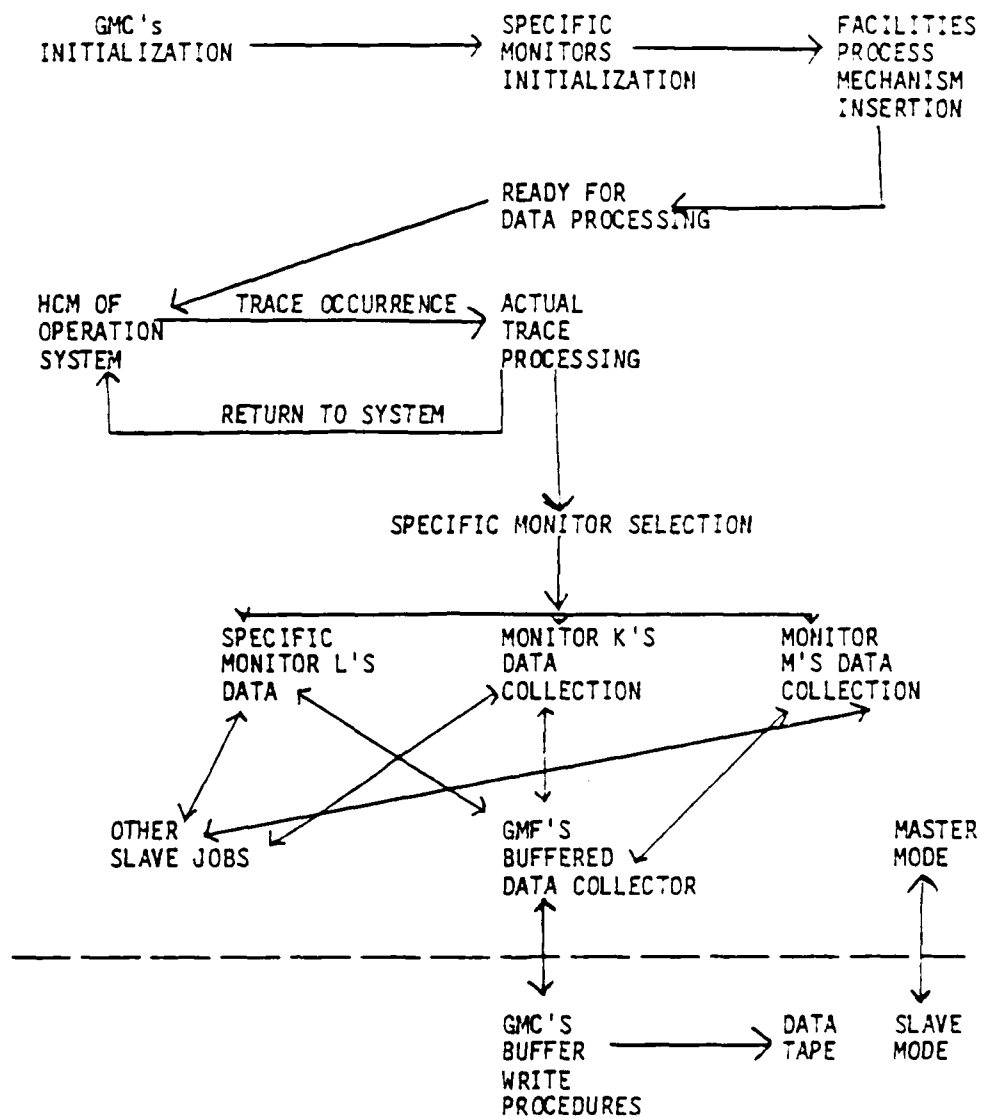


Figure 5-1. GMC Concept

After obtaining control, the system trace recording is completed using normal system procedures and then the trace is investigated. If an executing GMF monitor requires this trace type, control is given to that monitor. The monitor then collects data of interest to itself and requests the GMC executive to buffer the data. When the monitor completes its task, it returns control to the GMC executive. The GMC executive then transfers control back to the system trace processing routine within the GCOS dispatcher. The activities mentioned above take place in master mode under the guise of an extension of normal system trace procedures.

The second method used by GMC to gain control is for GMC to create its own system traces. The GMC will search a given GCOS module for a known line of code. It will replace this line of code with a transfer to a patch area. In the patch area, the monitor will insert code to create a new GMC system trace. At this point, the execution of this code will be processed just as all other system traces are processed. This procedure is used only when the Communication Analysis Monitor is selected for execution. (See subsection 5.2.6 for a complete description of this procedure.)

The third method used by GMC to gain control is via a direct transfer from a GCOS module. In this case, GMC will search a given GCOS module for a known line of code. It will replace this line of code with a direct transfer into the GMC. This can be accomplished since GMC is locked in core during its entire execution. The benefit of this procedure is that the overhead of the system trace is eliminated. This procedure is used only when the Mass Store Monitor or the CPU Monitor is selected for execution. (See subsections 5.2.2 and 5.2.3 for a complete description of this procedure.)

As the GMC ER buffers the collected monitor data, it will determine when one of the internal buffers are filled and must be written to tape. At this point, it will establish a normal slave dispatch to its tape writing facility. The tape writing facility will then write the internal buffer to tape and signal the executive that this buffer may be reused.

5.2 GMC Monitor Subroutines

In this subsection, each of the ten monitor subsystems will be addressed. Each subsystem requires that specific trace types be enabled in the H6000 system boot deck on the \$ TRACE card. A detailed discussion of the computer system boot deck used for startup and \$ TRACE operations is contained in the GCOS System Startup Manual and in the System Tables Manual. GMC cannot be used to turn on or off a TRACE option. The GMC user must request the computer system manager to change the system boot deck \$ TRACE card to meet the minimum GMC requirements. If all the required trace types are not on, GMC will abort with a TO through TS and TA abort. The hexadecimal digit immediately following the letter T indicates the monitor number for which the proper traces are not active. See table

5-1 for a quick reference of required trace types for each monitor and refer to table 5-2 for all GMC abort codes.

5.2.1 Memory Utilization Monitor. The Memory Utilization Monitor (MUM) is used to measure memory utilization. For a detailed description of reports containing data collected by this monitor, see section 6.

When MUM is active it is essential that GCOS system trace types (octal) 10, 11, 46 and 51 are enabled in the computer system boot deck \$ TRACE card. MUM collects data upon the occurrence of those traces and builds its records which are then passed to the Executive Routine (ER) for buffering. A separate discussion of the format of the MUM collected data records is contained in subsection 5.4.2. MUM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 2, 11, 19, 23, 26 and 27. The complete process for generating an R* file is described in subsection 5.6. It should be noted that this version of MUM will not collect any idle trace information and therefore not all MUM reports will be produced during data reduction. See section 6 for description of reports that will not be produced.

If the user wants all the MUM reports produced, then the Idle Monitor must be made active, along with the Memory Monitor. To do this, GCOS trace types (octal) 0, 1, 2, 3, 10, 11, 13, 16, 21, 22, 37, 46, 51 and 65 must be enabled. In generating the R* file the following segment numbers should be used: 1, 2, 10, 11, 19, 23, 24, 25, 26 and 27. It should be noted that when the Idle Monitor is active the amount of data collected will be about double that collected when the Idle Monitor is off. The user should carefully evaluate the need for those reports produced by the Idle Monitor.

The MUM is designed so that it can accurately report all changes to the memory subsystem. It does this by processing all trace type 10's (.CALL) and all trace type 11's (.GO TO). Upon receiving these traces a further check is made to see whether a memory allocation process is being requested i.e. the use of SSA modules: .MPOP3, .MPOQ4, or MPOR5. The MUM will collect data only if one of these modules has been requested.

The first item of information reported by the MUM is the current status of all jobs waiting in the Peripheral Allocator's Queue. This information is reported so as to be better able to represent the true core demand being made by the current workload. When a GCOS system has a large number of jobs waiting core, a Core Damper Switch is set. This switch is used to prevent jobs from being sent from the Peripheral Allocator to the Core Allocator. Therefore, the Peripheral Allocator's queue may contain many jobs that would normally be in the Core Allocator's queue, were it not for the Core Damper Switch. This

Table 5-1. Required Trace Type for Each Monitor

<u>Monitor #</u>	<u>Monitor</u>	<u>Required Trace Type</u>
0	Memory Utilization Monitor (MUM)	10, 11, 46, 51, (Idle Monitor traces optional)
1	Mass Storage Monitor (MSM)	7, 15, 73*, 76*
2	CPU Monitor (CPUM)	10, 11, 21, 70*
3	Tape Monitor (TM)	50, 51, 52
4	Channel Monitor (CM)	4, 7, 15, 22 (Idle Monitor traces optional)
5	Communications Analysis Monitor (CAM)	14*, 15
6	GRTS Monitor (GRTM)	62*
7	Transaction Proc- essing System Monitor (TPSM)	0, 1, 2, 4, 5, 6, 13, 42, 51, 65, 74*
8	Idle Monitor (IDLEM)	0, 1, 2, 3, 13, 16, 21, 22, 37, 65
A	TSS Monitor (TSSM)	74*

*These are not standard traces. They are specially created by the GMC or by an editing of the GCOS TPE Subsystem in the case of trace type 74. Trace types 70, 73 and 76 are direct transfers into the GMC and as such are not required to be active via the \$ TRACE card in the system boot deck. Trace types 14, 62 and 74 do use the System Trace Function and require the Trace Number to be active on the \$ TRACE card.

Table 5-2. Abort Codes (Part 1 of 3)

- |B2 - Illegal SNUMB on MSM data card (more than 5 characters).
- |B3 - More than 5 SNUMBs for MSM SNUMB option.
- BC - Illegal request on limited connect option.
- BK - Backspace of the full data tape was bad. Multireel will not be collected. Check for tape drive problems.
- BS - Bad tape status. Check condition of tape and rerun job.
- C1 - CPU Monitor turned off but SNUMB input requested on the data card.
- C2 - Illegal SNUMB (more than five characters) on data card for CPU SNUMB option.
- |C3 - More than three SNUMBs for CPU Monitor on data card.
- CD - Illegal character in CAM special option.
- CE - Console message garbled. Check console sheet and check with operator.
- CM - Cannot move out of the growth range of TSS.
- CO - CAM turned off but special option requested.
- DK - No multireel disk file was present. Use a \$ FILE card in the JCL or use the M9 option to turn off multireel capability.
- DR - Disk read-in. End-of-reel processing was bad.
- DS - Bad status of the disk write.
- ER - Wrapup record could not be written.
- ET - More than two terminals requested for CAM special option.
- |FN - The IOS accounting modification could not be found. Call CCTC.
- GC - No GRTS control card.
- GD - No FEP I/O can be performed.
- GM - Needed memory for GRTS Monitor denied. Increase \$LIMIT card.
- GO - GRTS Monitor illegal data card.
- |GS - Extra SSA is not available for GRTS Monitor. Check \$ LIMIT card.

Table 5-2. (Part 2 of 3)

- MO-M8,MA - Monitor was not turned off and not present on the R* file. Any monitor not contained on the R* file must be turned off via the data card option. The number following the M is the monitor that was not turned off.
- MM - The dispatcher hook has already been inserted. Another version of GMC must already be in execution.
- N1 - The CPU Monitor hook code could not be found. See subsection 5.2.3.
- N2 - Sufficient patch space is not available in .MDISP to run the CPU Monitor. See subsection 5.2.3.
- N3 - DNWW/MDNET patch location could not be found. See subsection 5.2.6.
- N4 - Sufficient patch space is not available in DNWW/MDNET to run the Communications Analysis Monitor. See subsection 5.2.6.
- N5 - MSM patch for SSA cache analysis not found (AOS .CRTDL). See subsection 5.2.2.
- N6 - MSM patch for SSA cache analysis not found (AOS .CRTBH). See subsection 5.2.2.
- N7 - MSM patch space in .MDISP not sufficient for monitoring SSA cache. See subsection 5.2.2.
- N8 - CPU Monitor hook code for subdispatch could not be found. See subsection 5.2.3.
- NF - The Dispatcher hook code could not be found. Call CCTC/C751.
- NS - A CAM abort because it could not find NSIP (# of special interrupts) address in .MDNET.
- NR - A CAM abort because it could not find ROLXCT (number of lines found waiting mailbox) instruction.
- OE - An error in an off option was encountered. Check the data cards. There is either an illegal character on the data card or a monitor which was not compiled in the R* file (see assembly listing) has not been turned off.
- OK - All went correctly.
- OL - Overlap of disk file. Increase size of disk file. Check if operator failed to respond to tape mount message during multiprocessing.

Table 5-2. (Part 3 of 3)

- OV - A tally overflow occurred in the MUM.T10 subroutine. Increase the size of the data area within subroutine MUM.T10, variable SIZEBUF.
- RS - Routine depth requested exceeded table length.
- RW - Error in initial rewind. Check tape and drive.
- SB - End-of-reel processing was bad. Check tape and drive.
- SD - Error setting of density.
- SF - Limited reel option termed successfully.
- SQ - Sequence error in the physical records.
- S1 - Subroutine MUM.T10 missing
- S2 - Subroutine MUM.T46 missing
- S3 - Subroutine CM.T07A missing
- S4 - Subroutine CPU.T70 missing
- S5 - Subroutine CM.T04A missing
- S6 - Subroutine CM.T22A missing
- S7 - Subroutine TM.T50 missing
- S8 - Subroutine CAM.T14 missing
- S9 - Subroutine GRT.T62 missing
- SA - Subroutine IDL.TRCS missing
- SC - Subroutine IDL.T21 missing
- SD - Subroutine TPE200 missing
- TE - The start/stop times appear improper. Check data card.
- TL - Trailer record write was bad. Check tape and drive.
- TS - An OK abort directed by a time to stop command.
- TW - The tally word has been garbled.
- ! TO-TB,TA - Required system trace is not on. The number following the T indicates the monitor having the problem.

Table 5-3. GMC Catalog and File Index (Part 1 of 4)

SEGMENT #	FILE	REQUIRED	FUNCTION
1	GMF.TOP	Y	Read data card, initialization, find hook in dispatcher, and create initial record
2	MUM.INIT		Initialize Memory Monitor
3	MSM.INIT		Initialize Mass Store Monitor
4	CPU.INIT		Initialize CPU Monitor
5	CAM.INIT		Initialize Communications Analysis Monitor
6	CM.INIT		Initialize Channel Monitor
7	TM.INIT		Initialize Tape Monitor
8	GRT.INIT		Initialize DN-355 Monitor
9	TP.INIT		Initialize TPE Monitor
10	IDL.INIT		Initialize Idle Monitor
10A	TSS.INIT		Initialize Timesharing Monitor
11	GMF.MID	Y	Ensure at least one active monitor
12	CAM.PAT		Preparation for patching DNWW/MDNET for Communications Analysis Monitor
13	CPU.PAT		Preparation for patching dispatcher for CPU Monitor
14	MSM.PAT		Preparation for patching dispatcher for MSM Cache Analysis
15	PATLOOK		Searches for patch space for CPUM, CAM, MSM
16	CPUDOIT		Patch the dispatcher for CPU Monitor
17	CAMDOIT		Patch DNWW for Communications Analysis Monitor
18	MSMDOIT		Patch dispatcher for MSM Cache Analysis

Table 5-3. (Part 2 of 4)

<u>SEGMENT #</u>	<u>FILE</u>	<u>REQUIRED</u>	<u>FUNCTION</u>
19	GMF.MON	Y	Insert the trace hook for GMC traces
20	CPU.REMO		Remove CPU Patches to dispatcher
21	CAM.REMO		Remove CAM patches to DNWW/MDNET
22	MSM.REMO		Remove MSM patches to dispatcher
23	GMF.BTM	Y	Remove the trace hook, do all IO control
24	IDL.TRC		Processes traces, 0,1,2,3,13,16,22,37,65 for Idle Monitor
25	IDL.T21		Processes trace 21 for Idle Monitor
26	MUM.T10		Processes traces 10,11,51 for Memory Monitor
27	MUM.T46		Processes trace 46 for Memory Monitor
28	CPU.T70		Processes traces, 10,11,21,70 for CPU Monitor**
29	TM.T50		Processes traces 50,51,52 for Tape Monitor
30	CAM.T14		Processes traces 14 and 15 for CAM*
31	CM.T04A		Processes trace 4 for Channel Monitor
32	CM.T22A		Processes trace 22 for Channel Monitor
33	CM.T07A		Processes traces 7,15,73,76 for Channel Monitor and Mass Store Monitor**
34	GRT.T62		Processes trace 62 for GRTS Monitor*
35	GRT.COL		Interfaces with the DN-355

Table 5-3. (Part 3 of 4)

<u>SEGMENT</u> <u>#</u>	<u>FILE</u>	<u>REQUIRED</u>	<u>FUNCTION</u>
36	TPE200		Processes traces 0,1,2,4,5,6,13,42,51,65 and 74 for TPS Monitor*
36A	TSS.COL		Captures trace 74 for TSS Monitor*
37	RUN.GMF		JCL to run a GMC R *
38	GMF.OBJ		File to contain a GMC R *
39	MAKE.XXX		A series of files creating different GMC R * Monitors.
39A	MAKE.MUC		Memory and CPU Monitors
39B	MAKE.ALL		Total GMC
39C	MAKE.MUM		Memory Monitor
39D	MAKE.CPU		CPU Monitor
39E	MAKE.TM		Tape Monitor
39F	MAKE.MSM		Mass Store Monitor
39G	MAKE.MCC		Memory, CPU, Communications and Idle Monitors
39H	MAKE.MCI		Mass Store, Channel and Idle Monitors
39I	MAKE.CAM		Communications Analysis Monitor
39J	MAKE.CM		Channel Monitor
39K	MAKE.GRT		DATANET-355 Monitor

Table 5-3. (Part 4 of 4)

<u>SEGMENT</u> <u>#</u>	<u>FILE</u>	<u>REQUIRED</u>	<u>FUNCTION</u>
39L	MAKE.CMI		Channel and Idle Monitors
39M	MAKE.MCM		Mass Store and Channel Monitors
39N	MAKE.GC		Communications and DATANET Monitors
39O	MAKE.TPE		TPE Monitor
39P	MAKE.TSS		TSS Monitor

*Trace types 14,62 and 74, are not standard. They are internally generated (IT) traces.

**Trace types 70,73 and 76 are not standard. They are direct transfer (DT) traces.

THIS PAGE INTENTIONALLY LEFT BLANK

5-11.2

CH-3

information from the Peripheral Allocator is reported only when the Peripheral Allocator is in memory and a Memory Monitor trace is about to be generated. For this reason, not all Peripheral Allocator queue changes will be reported. In order to reduce the amount of information being collected, a job's status in the Peripheral Allocator's queue is reported only for new jobs, when a job has changed activity, or when its status has changed.

After reporting any Peripheral Allocator status information, the MUM will next report the status of every job waiting for or currently using memory. Information such as the SNUMB, IDENT, USERID, Activity Number, memory demands, current memory address, whether the job is in memory or waiting for memory, and whether the job is a system program or user program is collected. This information is reported for each job only if a change has occurred from previous information that was reported. In addition, the current amount of CPU and IO time used by a job is reported in every MUM trace that is generated.

The MUM will consider a job to be a system job if it has a program number less than octal 10, or if it has no J* file and requires privacy. Since the user may want to consider other jobs to be system jobs, such as HEALS or VIDEO, the data reduction program allows the user to extend this definition of system jobs (see section 6).

5.2.2 Mass Storage Monitor. The Mass Storage Monitor (MSM) is used to collect data on usage of peripheral resources. For a detailed description of reports containing data collected by this monitor, see section 7.

When the user wants MSM to be active, it is essential that trace types (octal) 7 and 15 are enabled in the computer system boot deck on the \$ TRACE card. MSM processes trace types 7, 15, 73, and 76 to build its own records which are passed to ER. A separate discussion of the format of the MSM collected data records is contained in subsection 5.4.3. As has been stated earlier, trace types 73 and 76 are direct transfer traces created by the GMC, and as such need not be defined on the \$ TRACE card. The MSM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 3, 11, 14, 15, 18, 19, 22, 23, and 33. The complete process for generating an R* file is described in subsection 5.6. If the system being monitored by the Mass Store Monitor contains SSA Cache Core, two new direct transfer traces, are created by the Mass Store Monitor in order to collect sufficient data to be able to analyze the operation of SSA Cache Core. These traces are created only if SSA Cache Core is configured. The Mass Store Monitor searches the dispatcher for a AOS .CRTDL instruction and then inserts code to make a direct transfer into the GMF. In addition, an AOS .CRTBH instruction is also searched for so that another direct transfer into the GMC can be generated. The first instruction locates the area of the dispatcher where it has been determined that the required SSA module is not in the SSA Cache Buffer and needs to be loaded from disk. The second instruction

and 2460. If GMC cannot find the ASA .SALT,5 instruction, it will abort with an N1 abort; if it cannot find the STQ instruction it will abort with an N8 abort. If either abort occurs, the dispatcher code should be examined to determine if either instruction has been modified, moved, or patched. If so, the code in CMC will need to be modified.

Upon finding these instructions, GMC searches the dispatcher patch area(s) for four free locations under WW6.4 or eight free locations under WW7.2 in which to create a direct transfer trace into the GMC. This search has the same ranges as that for SSA cache in MSM. If patch space is not found, an N2 abort will occur. See subsection 5.2.2 for a description of this search procedure.

The CPU Monitor tracks the CPU usage of all system programs and accumulates CPU usage of slave jobs into a single value (see subsection 5.4.4). If the user desires, he can specify up to three slave jobs for which he wants the CPU monitor to track CPU usage, just as it does for system jobs. Subsection 5.5.5. describes this user option.

5.2.4 Tape Monitor. The Tape Monitor (TM) is used to collect utilization statistics on magnetic tape drive activity. A separate discussion of the format of the TM collected data records is contained in subsection 5.4.5. Reports containing data collected by this monitor are described in section 11.

When the user desires that the TM be active, GCOS trace types (octal) 50, 51, and 52 should be enabled in the computer system boot deck on the \$ TRACE card. TM processes these trace types to build its records which are passed to the ER. The TM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 7, 11, 19, 23, and 29. The complete process for generating an R* file is described in subsection 5.6.

5.2.5 Channel Monitor. The Channel Monitor (CM) is used to measure I/O channel activity over tape and disk channels and contention to disk devices. A separate discussion of the format of the CM collected data records is contained in subsection 5.4.6. See section 8 for a description of reports containing data collected by this monitor.

When CM is active, it is essential that GCOS trace types (octal) 4, 7, 15, and 22 are enabled in the computer system boot deck on the \$ TRACE card. CM processes these trace types to build its records, which are passed to the ER. The CM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 6, 11, 19, 23, 31, 32, and 33. The complete process for generating an R* file is described in subsection 5.6.

Actually, when the CM is active, sufficient data is processed for obtaining reports not only from the Channel Monitor but also from the Mass Store Monitor. The only Mass Store Monitor data that cannot be

collected would be the data needed to analyze Cache Memory. If the user also wants this data to be collected, he should create an R* file from the following segments (see table 5-3): 1, 3, 6, 11, 14, 15, 18, 19, 22, 23, 31, 32, and 33. In addition, the Mass Store Monitor must be made active. There is an additional option available with the Channel Monitor. This option allows the Channel Monitor Data Reduction Program to produce a CPU Idle/IO Active Report. This report is described in section 8. To obtain this report, the Idle Monitor must be included in the R* file. In addition, all Idle Monitor traces must be active. The following segments are required to generate the R* file: 1, 6, 10, 11, 19, 23, 24, 25, 31, 32, and 33.

5.2.6 Communications Analysis Monitor. The Communications Analysis Monitor (CAM) is used to measure machine and user response time and terminal usage. A separate discussion of the format of the CAM collected data records is contained in subsection 5.4.7. The complete process for generating an R* file is described in subsection 5.6. The output reports, containing data collected by CAM, are described in section 9. When CAM is active, it is essential that the GMC generated trace type (octal) 14 and the GCOS trace type (octal) 15 are enabled in the computer system boot deck on the \$ TRACE card. CAM patches the DNWW (MDNET in W7.2) module, looking for a LDQ M.LID,3 instruction followed by an ANQ =0077777,DU instruction. When the monitor is terminated, CAM removes these patches from the system. The CAM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 5, 11, 12, 15, 17, 19, 21, 23, and 30.

The method used by the CAM to patch DNWW/MDNET is similar to that used by the CPUM to patch the dispatcher. The CAM searches DNWW/MDNET for the LDQ M.LID,3 instruction beginning at octal location 5000 and ending at octal location 6000 (offsets from the entry point). If CAM cannot find this instruction, GMC will abort with an N3 abort. If this problem occurs, the DNWW/MDNET code should be examined to see if this instruction has been moved out of the octal range 5000-6000 due to an edit or recompile. If so, the code in CAM.PAT will need to be altered.

Upon finding this instruction, CAM then searches DNWW/MDNET patch area for 10 free locations in which to create a new system trace type 14.

| This search begins at octal location 11200 and continues for 100 octal locations (offsets from the entry point). If 10 free words of space are not found, then seven words of patch space are searched for within the dispatcher. This search occurs between octal locations 3540-3740 in W6.4 or 4600-5000 in W7.2 (offset from the entry point). If no space is found by either of these searches an N4 abort will occur. In this case, the user should examine the patch deck to see if a large number of patches have been made to DNWW/MDNET. If this is the case, DNWW/MDNET will need to be re-edited in order to remove these patches or else the CAM will not be able to be utilized. In addition to the above patching, CAM.INIT also searches DNWW/MDNET for certain

instructions. Beginning at 1400 octal locations from the entry point, and continuing for 5000 octal locations, CAM.INIT searches for a ANA=0777777,DL, followed by a CMPA instruction. If it does not find these instructions, it will abort with an NS abort. CAM.INIT is searching for a number of special interrupts processed (NSIP). In addition, CAM.INIT will also search for a ANA=077, followed by a CMPA=077 instruction beginning at the above CMPA location and continuing for 3000 octal locations. If it does not find this instruction, it will abort with an NR abort. In this section, CAM.INIT is searching for the ROLXCT processing (number of lines found waiting mailbox). If a specific terminal's traffic is to be monitored (see subsection 5.5.3), the CAM will insure that no more than two terminal IDs have been included. Invalid terminal IDs, extra terminal IDs or terminal IDs without the CAM input option request will cause the GMC to abort with a CD, CO, or ET abort. See table 5-2 for the meaning of these aborts.

The CAM also uses the GCOS trace type 15 (octal) to check for any JDAC processing or any other line switching which may occur.

5.2.7 GRTS Monitor. The purpose of the GRTS Monitor (GRTM) is to collect statistical data to be used in evaluating the performance of the DATANET 355 Front-End Processor (FEP). This data includes CPU Utilization, Response Time, and Terminal Performance. The collected information is sent to the GMC within the H6000, which writes the data to tape. A separate discussion of the format of the GRTM collected data records is contained in subsection 5.4.8. This tape, containing GRTM performance data and possibly data from other monitors, is used as input to a data reduction program used to produce statistical reports. (See section 10).

5.2.7.1 Configuration Requirements. The GRTM will execute on H6000 system with up to eight FEPs. The monitor is designed to run on the GRTS II Phase IIA (GRTS 1.2) FEP software.

5.2.7.2 H6000 Configuration Requirements. To run GRTM, GCOS trace type (octal) 62 must be enabled via the H6000 computer system boot deck on the \$ TRACE card. The GRTM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 8, 11, 19, 23, 34, and 35. The complete process for generating an R* file is described in subsection 5.6.

5.2.7.3 Altering of Phase II-A Software. To use the GRTM, the user must modify the standard GRTS software by applying a set of alters supplied with release of the GMC software. It should be noted that in Release WW7.2 the alter cards to support the monitor are included within the standard release. The user must check the FMAC module to insure that variable FMON has been set to 1. The FMAC module must be recompiled and the macro library reloaded. Finally, all the GRTS modules should be recompiled.

B29IDPX0		<u>GMFCN</u>							
NAK.XXX	GMF.OBJ	PATLOOK	GMF	MUM	MSM	CM	CPU	TM	
			GMF.TOP	MUM.INIT	MSM.INIT	CM.INIT	CPU.INIT	TM.INIT	
			GMF.BTH	MUM.T10	MSM.PAT	CM.T04A	CPU.PAT	TM.T50	
			GMF.MID	MUM.T46	MSM.REMO	CM.T22A	CPU.T70		
			GMF.MON		MSMD01T	CM.T07A	CPU.REMO		
							CPUD01T		

Figure 5-2. GMC Catalog File Structure (Part 1 of 2)

As an example, a PTA value of 317200 would be used to identify the pseudo terminal as being on:

```
IOM CHANNEL NUMBER OF HSLA = 6
HSLA NUMBER                 = 1
HSLA SUBCHANNEL             = 29
POLLED SCREEN NUMBER        = 0
MUST BE ZERO
```

This is the same format used to describe subchannels within .MSECR.

A user must insure the setting of the PTA value as circumstances demand. By using the format shown above, the symbolic location PTA located in the FSUB module may be altered to reflect the user's own PTA value. The FSUB module must then be reassembled prior to bootloading the DATANET.

5.2.7.6 Abort Codes. The general abort codes listed in table 5-2 apply to specific options. Abort codes created specifically for the GRTM are listed below.

```
GS = SSA processing error caused by missing or invalid $LIMITS
    card

GC = Invalid GRTM options on data card file (I*)

GC = Missing GRTM option card

GD = No FEP I/O can be performed

GM = GEMORE unsuccessful in getting needed buffers
```

NOTE: GM aborts occur if another program occupies continuous memory just above that of the GRTM when buffer space is requested using MME GEMORE. Increasing the \$LIMITS memory value or loading the GRTM immediately after booting the system will enhance the chances of getting DATANET buffers.

5.2.7.7 DATANET Monitor Software Description. The GRTS II system operating within the DATANET will be used in the collection and recording of various internal resource information which is then sent to the GMC executing in the host (6000) system. Monitoring functions within the DATANET have been separated into three basic monitoring categories:

- a. CPU and Resource Monitoring
- b. Host/FEP Response Time
- c. HSLA Subchannel Monitoring

Monitor information will be passed between the DATANET and the GMC program executing in the host using the normal FICM interface.

5.2.7.7.1 DATANET-HOST Interface. The GRTM executing in the DATANET will be in the form of a pseudoterminal connected to the DATANET. Once initialized, the pseudomonitor terminal will be as any other remote terminal connected to a Direct Access (DAC) program in the host.

As part of GRTM initialization, the pseudo terminal will issue a connect-to-slave request to the DAC GMC program in the host. The DAC connect name requested will be the special monitor name of GRTMN'X' with the 'X' being a number from zero through seven which corresponds to the DATANET (0-7) that issues the connect request.

The connect-to-slave request will remain outstanding until it is answered by an inquiry issued by the GMC program. Once the request has been answered, the GRTM/GMC program connection will be made.

Since the pseudo monitor terminal will not be physically connected to the DATANET (i.e., on an HSLA S/C) the need for a special monitor "Terminal Control Block (TCB)" becomes necessary.

The special monitor TCB is necessary in order to utilize the normal FICM interface in the passing of monitor information to the GMC program.

5.2.7.7.2 Monitoring of CPU. The DATANET Monitor will collect CPU and various resource information by the placement of conditionally assembled instructions at appropriate points in the ICM and EXC modules of the GRTS II software. The information to be collected at the CPU level is independent of the HSLA subchannel and is sent to the host based upon a predefined time sampling increment for writing buffers to the host collector program. For a given buffer sent to the host, the following information is included:

- a. Time Idle - The amount of time in milliseconds spent in the exec idle loop since the last buffer was sent to the host.
- b. Buffer Denials - A cumulative count of buffer denials. This is a count of the number of times that the GRTS II software was unable to get a buffer for a user.
- c. Buffer Availability - The number of 18-bit words currently available for buffers.
- d. Number of Users - Count of the number of users currently logged on the system.

System. A separate discussion of the format of the TPSP collected data records is contained in subsection 5.4.10. The reports containing data collected by TPSP are described in section 12.

When TPSP is active, the required traces must be enabled in the computer system boot deck on the \$ TRACE card (see table 5-1). A sample of the reports and run time procedures for the data reduction program can be found in the Transaction Processor System section 12. The TPSP requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 9, 11, 19, 23, and 36. The complete process for generating an R* file is described in subsection 5.6.

NOTE: The TPSP cannot be run concurrently with the TSSM.

5.2.9.1 TPS Trace Collection. The TPSP is unlike most other GMC monitors in that monitoring of the Transaction Processing System is controlled via the operator console. Prior to collecting data, the user must alter the TPS (see subsection 5.2.9.2) and must also create a usable GMC R* file (see subsection 5.6). Once these actions are performed and a GMC execution is started, the user must still perform one additional action before data collection can begin. The TPSP is turned on or off by the console operator via the TP MESS command. The operator must request "TP MESS". When the console responds "TP MESS?", the message "TRACE ON" is entered to start processing traces, or "TRACE OFF" to discontinue trace processing. This procedure can be repeated as often as desired. The TPSP and the TSSM are the only GMC monitors that can be turned on or off while the GMC is physically executing.

5.2.9.2 Modifying the Transaction Processing System. To use the TPSP, the user must alter the Transaction Processing System. An alter file is provided with the GMF software delivery. The file name is B29IDPXO/SOURCE/TPEALT. This file contains all alters and associated JCL necessary to modify the standard TPS. These modifications apply only to the WW7.2 version of TPS. The user will need to make minor modifications to the file so as to correctly reference any permanent files that are required.

5.2.10 Timesharing Subsystem Monitor. The Timesharing Subsystem Monitor (TSSM) is used to measure TSS performance. Section 15 details those reports available from the data collected by this monitor.

When TSSM is active, GCOS trace type 74 (octal) must be enabled on the boot deck \$ TRACE card. The TSSM causes the trace to be taken from many points in TS1; the collector builds its records which are then passed to the ER for buffering. An example of the record format appears in subsection 5.4.11. TSSM requires the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 10a, 11, 19, 23 and 36a. Subsection 5.6 shows how to generate an R* file.

If all TSSM reports are wanted, then both the CPU Monitor and Channel Monitor are additionally required: GCOS trace types (octal) 4, 7, 10, 11, 15, 21 and 22 must also be enabled. Restricting Channel Monitor activity to SNUMB TSI will conserve tape and GMC overhead. The TSSM and TPSM cannot be run concurrently.

The TSS Monitor consists of two major pieces of software: a data generation program loaded into TSS and a data capture routine loaded along with GMC routines. The TSS Monitor is used to analyze periods of poor TSS response to determine which users are affected, when and for how long they are affected, and what the possible causes of the poor response are. To collect information for such analysis, probes are inserted throughout TSS to gather individual pieces of information which, when combined by a data reduction program, yield the desired result. When active, the TSS Monitor is an area of code entered via transfer instructions planted throughout TSS by an initialization routine. After the TSS Monitor is entered from a probe point, it makes an entry in the GCOS trace table and then returns to the TSS process which was executing. The places where transfers are made fall into the following categories:

- o Session identification and correlation - retrieval of line ID, DRUN, ID, USERID and User Status Table (UST) address (the unique identifier which is used to correlate traces coming from an individual session)
- o Type of user interaction - retrieval of subsystem names, indication of build mode, memory sizes, derail codes
- o Terminal I/O - a single location where I/O is started and a number of locations where courtesy calls for I/O complete are paid
- o Disk I/O to user files - a single location where I/O is started and a number of locations where courtesy calls are paid for I/O complete
- o FMS service - pairs of locations, one where TSS issues a request for service, and the other where the service is complete; the second location also gives an FMS status to tell whether the service was actually performed or why it was denied
- o Processor allocation - locations for entering the processor allocation process, placing an entry into the subdispatch ready queue, and removing an entry from the subdispatch fault queue
- o Memory allocation - retrieval of subsystem size; recording of the flow of control in the memory allocation and swap processes
- o Errors and denials - retrieval of numeric error codes, recording of individual denial events

origin to be at location 3760 with respect to the beginning of TSS0. The first patch NOP's out a conditional transfer instruction around an instruction which loads a UST origin corresponding to the case when VIPs are configured. The letter "O" in these patches means that the patch locations are with respect to the origin of TSS0, not to slave address 0 of TSS as in the subsystem attribute patches. The letter "R" before patch content indicates that the address field of the patch must be relocated to the beginning of the module (TSS0) identified in column 7 of the patch. The symbolic addresses for these patches are IN7630+2 and IN7630+3. The following patches complete the TSS INIT file when the TSS Monitor is loaded on a startup file:

5271	OCTAL	11007	DON'T TRANSFER IF NO VIPs	.MTIMS
5272	OCTAL	R3760235007	ALLOW 1K FOR GMF	.MTIMS

5.2.10.3.3.4 Installation from a Permanent File. The principle of this method is that if a job has a file code ** active, any MME GECALL will cause that file to be searched before the system files are searched. Patches on the TSS INIT file must access a permanent file before subsystem loading begins and release it after subsystem loading finishes. The TSS User Derail Loader (TUDL, SDN K79005) provides an example of how to access and release a system loadable file from within the TSS executive. The patches on the TSS INIT file for the TSS Monitor are compatible with TUDL because TUDL starts its work after the TSS Monitor has finished its work. TUDL may further relocate the UST origin upward, but TUDL uses the address stored by the TSS Monitor patches. A list of an entire \$ PATCH section is given in the TSSM source code. The first four patches match ones described in earlier sections.

The bulk of the patches used to access a permanent file are placed in an area of TSS0 which is reserved for UST space when the UST origin is not adjusted upward (actually, this space was needed in releases prior to W7.2.0 to prevent the generation of a UST for TSRL from destroying the code at the end of TSS startup; with the TSS INIT file feature, enough code intervenes so that this buffer is not necessary and so that the UST origin can be moved up and not destroy the code). The first transfer into these patches is at offset 4673 in TSS0 (symbolic offset DMYERR+2, instruction EAX5 1). Here, the USERID in the file structure defined in the patches is stored into the SSA of TSS at location .SUID. Then a MME GEMORE accesses the file. Use of .SUID makes FMS think that the file is being accessed by its owner and thus, no special permissions are needed. If the MME GEMORE is denied, a flag is set to 0 before control returns to the original TSS0 coding from offset 2013 in the patches.

| The second transfer, at octal offset 5272 (symbolic location IN7630+3) replaces the instruction defining the UST origin. The patch at octal offset 5271 remains intact. In the second group of patches, the

permanent file must be released, the number of PATs decremented by one, and the USERID in .SUID set to zero. If the GEMORE in the first set of patches is denied, then the number of PATs is not decremented. This alteration of word .SNPAT in the SSA is necessary because releasing the file does not remove the file code from the SSA. The patch for defining the new UST origin is moved to octal offset 2031 in the patches, just before a transfer back to TSS0.

5.2.10.4 Production Use of the "Monitor. The following steps must be taken to enable data capture:

1. Append patches on file B29IDPX0/GMFCOL/TSS/TSS.PAT to Timesharing INIT File (normally OPNSUTIL/TS1).
2. Run the file B29IDPX0/GMFCOL/TSS/TSGMF to create the new TSS subsystem file.
3. Start TS1.
4. Start a copy of GMC with the TSS Monitor active.
5. Log on to a master USERID and enter "SYST GMF" at the "*" prompt following log-on to install the hooks into TSS code. Traces will not start at this time. The master USERIDs assembled in TSS are MASA and MASB. Unless these are patched or redefined in the TSS INIT file, only a terminal ID designated as master in .MSECR may be used for this step.
6. Enter "TS1 TRACE ON" from the system console to start generation of traces from TSS.
7. Enter "TS1 TRACE OFF" from the system console to suspend generation of traces.

Steps (5), (6) and (7) execute independently of (4); however, use of step (6) without use of step (4) will cause unnecessary TSS overhead if traces are being generated and lost due to GMC not being in execution. Steps (6) and (7) may be repeated multiple times if traces are to be captured for specific periods of the day. (The companion data reduction program, described in section 15 cannot process GMC sessions longer than 9 hours).

5.2.10.5 Monitor Limitations. To obviate lockup fault, initial trace generation (at "TS1 TRACE ON", and following lost data) is inhibited until the number of TSS users falls below 50. Further, all trace generation is inhibited until the number of users falls below 100. (The companion data reduction program is limited, by parameter, to 50 active USTs).

Whenever the dispatcher executes a trace after this point, the GMC ER will gain control and pass the system control over to the proper monitor for data collection. Control is then returned to the ER for return to the dispatcher.

In the process of collecting data, a monitor may desire to save a logical record on the collection tape. In order to sequence all logical records to the tape file, a common buffering system has been provided by the GMC. Any monitor can pass control to symbol BUFCTL for buffering a logical record to be saved on the tape file. This code is executed in master mode as an extension of the operating system function where control has been taken. Writing the buffered data to tape is a GMC slave function which requires controls between the buffering system and the writing system. The buffering system indicates, via software to the slave, when an individual buffer is to be written. After the slave has written the buffer, it is returned to the buffering system. The slave code for writing the tape is started at symbol BUFCHK. This slave portion is normally in GEWAKE mode and is awakened either by GMC master mode code or by a normal dispatch. The slave function will check for a full buffer and write it before going back to sleep. This process is repeated until GMC has been terminated, at which time the slave portion will execute the proper abort code via a MME GEBORT. The wrap-up procedure then performs all termination requirements. GMC will terminate via direction from the time option on the data card, limited reel option on the data card or else via console command. If no time option or limited reel option has been selected, the GMC will continue to process until it is ABORTED via the console.

At termination, the wrap-up location is given control and the termination record is processed. This is the slave termination of GMC. Also associated with termination is the return of the dispatcher code to its original state. This is accomplished by GMC checking its own .STATE word for a termination code. When the termination code is sensed, the dispatcher hook is removed, the original code is replaced, and the slave termination process is completed. Upon termination, ER reads the communication region table, writes the data to tape with a termination record, removes any patches it has placed in the system, and proceeds to wrap up.

5.3.2 Output. Output of GMC is always a magnetic tape file. The following is a guideline for the number of tapes generated by each monitor (800 BPI, 2400 feet tape reels):

- MUM - one tape every 6-8 hours. With Idle Monitor, one tape every 3 hours.
- MSM - one tape every hour

CPUM - 1 tape every 24 hours
 TM - 1 tape every 24 hours
 CM - 1 tape every half hour
 TPSM - 1 tape every 4 hours
 CAM - 1 tape every 8 hours
 GRTM - 1 tape every 24 hours
 TSSM - 1 tape every 2 hours

Analysis of this tape(s) is accomplished by a series of dedicated data reduction programs which are discussed in sections 6 through 12 and section 15. It is essential that the proper data reduction program is run against data created by its associated monitor routine.

The GMC is usually terminated by operator request. However, there are times when the monitor may terminate itself. In most of these cases, the operator will receive a console message indicating the reason for the abort (see table 5-2). However, under some circumstances the message the operator sees is that GMC terminated with "NO REASON SPECIFIED." In this case, if the job listing is examined, a reason will be given under the WRAPUP activity information as shown below.

* ACTY-01 \$ CARD #0008* GELOAD 03/16/78 SW=000000000000
 NO REASON SPECIFIED AT 030440 I=5060 SW=000000000000

*WRAPUP BEGUN
 (USERS BS MME GEBORT) AT 031413 I=0000 SW=000000000000

5.4 GMC Data Records

5.4.1 GMC Executive. The GMC Executive produces an initialization record that must be read by every data reduction program. This record contains information on the system configuration and the status of various queues at the time the monitor started. The length of the record is dependent on the size of the configuration. The Executive also writes a termination record whenever it terminates normally.

Initialization Record

<u>Word</u>	<u>Bits</u>	<u>Information</u>
1	0-35	Block Control
2	0-35	Zero
3	0-17	Year (.CRJCD)
	18-35	Julian day (.CRJCD)
4	Not used	
5	0-35	Current date (.CRDAT)
6	0-35	Current time (.CRTOD)
7	0-35	Reg A of RSCR 32

CC	
1	
MO M3	Turn off Monitor 0 and 3
M1	Turn off Monitor 1
M1 M9	Turn off Monitor 1 and collect only a single reel
M1 *12.36,05.00	Turn off Monitor 1, start collecting data at 12.36, and collect for 5 hours
.03.00	All Monitors are present on the R file and are active, collection is to start at once and continue for 3 hours
+CK	All Monitors are present on the R* file, and communication traffic is to be monitored for terminal CK
M1 M4 M93	Turn off Monitors 1 and 4, and collect maximum of three reels of data
M*	Suppress abort if GMC cannot move
#VIDEO,HEALS	All Monitors are present on the R* file, and accumulate processor time in the CPU Monitor for these SNUMBs.
MO M5 M8 ?1	Turn off monitors 0, 5, and 8. Collect only tape connects with the MSM and CM.
MO M2 M5 M; X MS2755TRTOS	Turn off monitors 0,2,5, set tape density to 1600 BPI, read second data card, turn on MSM/CM special SNUMB option to include TSS, FTS, \$PALC, 2755T and RTOS.
MO M4 M; MS	Turn off monitors 0,4, set tape density to 1600 BPI and collect MSM/CM traces for only TSS, FTS, \$PALC.

Figure 5-3. Data Card Examples

- (3) Requesting complete communication data for 1 or 2 terminal IDs.
- (4) Suppressing a GMC abort if it cannot move to an acceptable location.
- (5) Specifying up to three SNUMBS to be processed by the CPU Monitor.
- (6) Requesting that only tape connects or mass storage connects be collected, but not both. The default is to collect both types.
- (7) Declaring the start and stop times of monitoring.
- (8) Requesting high density tape be collected.
- (9) Specifying that the Mass Store Monitor and/or Channel Monitor are to collect data only for certain jobs.
- (10) Specifying that the data card options are continuing on a new card.
- (11) Specifying the monitoring requirements for the GRTM.

5.5.1 On/Off Option. This option allows the user to turn off all monitors not required for his purposes. Since the GMC default is to have all monitors turned on, unless specifically turned off, and since the TSS and TPE Monitors are incompatible, the user must have a data card and at least one of these two monitors must be turned off. The code format to turn off a given monitor is:

M0 = Memory Utilization
 M1 = Mass Store Monitor
 M2 = CPU Monitor
 M3 = Tape Monitor
 M4 = Channel Monitor
 M5 = Communications Monitor
 M6 = GRTS Monitor
 M7 = TPE Monitor
 M8 = Idle Monitor
 MA = TSS Monitor
 MB-MF = User Developed Monitors

(See Section 13 for a discussion of user developed monitors.)

CAUTION: The TPE Monitor and the TSS Monitor are incompatible and cannot be active at the same time.

Table 6-1. (Part 2 of 4)

<u>ID Number</u>	<u>Histogram Title</u>
21	Number of Activities Waiting Memory When a Processor Went Idle *
22	Memory Available When a Processor Went Idle *
23	Delay Time in the System Scheduler
24	Delay Time Until Core Allocation
25	Percent of Assigned Memory Used (Time-Corrected)
29	Number of User Activities Waiting Memory in the Allocator Queue *
30	Number of User Activites in Memory *
31	Elapsed Time of a Busy State Processor 0
32	Elapsed Time of a Busy State Processor 1
33	Elapsed Time of a Busy State Processor 2
34	Elapsed Time of a Busy State Processor 3
35	CP Time Per User Activity
36	I/O Time Per User Activity
40	Number of Activities Waiting Memory (Time-Corrected) *
41	Number of Activities in Memory (Time-Corrected) *
42	Memory Available (Time-Corrected) *
43	Number of User Activities Waiting Memory (Time-Corrected) *
44	Number of User Activities in Memory (Time-Corrected) *
45	Total Demand Outstanding (Time-Corrected) *

* Jobs with 0 urgency are not included.

Table 6-1. (Part 3 of 4)

<u>ID Number</u>	<u>Histogram Title</u>
46	Number of Extra Activites That Could Fit in Memory Without Compaction
48	Length of an Idle State (All Processors)
49	Length of an Idle State Processor 0
50	Length of an Idle State Processor 1
51	Length of an Idle State Processor 2
52	Length of an Idle State Processor 3
53	Number of Times System Activity Swapped
54	Elapsed Time a System Activity was Swapped
55	Elapsed Time of a Busy State Processor 4
56	Elapsed Time of a Busy State Processor 5
57	Length of an Idle State Processor 4
58	Length of an Idle State Processor 5
<u>ID Number/Name</u>	<u>Plot Title</u>
26/PLOT1	Availability of Memory vs. Outstanding Demand In Core Allocator Queue vs. Outstanding Demand in Peripheral Allocator Queue Plus Outstanding Demand in Core Allocator Queue
27/PLOT2	Memory Shortfall in Core Allocator Queue vs. Memory Shortfall in Core Allocator Queue Plus Memory Short- fall in Peripheral Allocator Queue
28/PLOT3	Number of Activities in Core Allocator Queue vs. Number of Activities in Peripheral Allocator Queue
59/PLOT4	Average size of TSS, FTS and NCP

cover a larger range of values. This change could be made via data cards and would not increase the size of the program.

The second method would involve increasing the size of the histogram by altering the value of TABSIZ. As long as the size requested does not exceed 50, this change can also be done via a data card. However, if an individual histogram needs to be larger than 50 buckets, the user will need to change the value of MXTBSZ. This change will require a change to source code, a recompile, and probably, an increase in program size. All references to MXTBSZ must be altered. This would need to be done in the EDIT subsystem of Time-Sharing.

The remaining items that can be modified are the title and the vertical axis headers. Table 6-2 shows the default values for all histograms.

6.1.4 Plot Options. There are three characteristics directly available to the user for each individual plot axis used.

The first characteristic, MAXNUM, is the maximum number of entries to be plotted on each vertical plot axis.

The second characteristic, YMAX, defines the upper limit of the horizontal display axis.

The third characteristic, YMIN, defines the lower limit of the horizontal display axis. Table 6-3 shows the default values for all plots.

6.1.5 Default Option Alteration. The general format for an option request is as follows: The first card contains an action code describing the action to be taken. Subsequent cards modify report parameters for some of the action codes. All input cards are free format with the only requirement being that at least one blank space separates multiple input parameters. The very last input card must have the word "END" typed on it. This card must be present whether or not any other input options are selected. Available actions with their (default) implications are shown in table 6-4. There is no order required for the options. In reading the following sections it should be remembered that the first card for any input option must be the action code specification with no other data present on the card.

The user should take special note that if this software is executed under a WW6.4/2H GCOS release, an additional data card is required. This data card is not described elsewhere in this chapter. The data card should contain the letters RN:

Table 6-2. Default Values for Histograms (Part 1 of 2)

<u>ID #</u>	<u>Low Value</u>	<u>Interval Size</u>	<u>Number of Buckets</u>
1	4	4	50
2	0	50	50
3	0	250	50
4	0	1	50
5	4	4	50
6	0	1	50
7	0	5	50
8	0	200	50
9	0	200	50
10	.95	.1	50
11	4	4	50
12	0	10	50
13	0	1	50
14	0	1	50
15	0	1	50
16	0.0	5.0	50
17	0	10	50
18	4	10	50
19	4	20	50
20	0	1	50
21	0	1	50
22	4	8	50
23	0	25	50
24	0	25	50
25	50	2	50
29	0	1	50
30	0	1	50
31	0.0	5.0	50
32	0.0	5.0	50
33	0.0	5.0	50
34	0.0	5.0	50
35	5	5	50
36	5	5	50
40	0	1	50
41	0	1	50
42	0	10	50
43	0	1	50
44	0	1	50
45	0	10	50
46	0	1	50
48	0.0	5.0	50
49	0.0	5.0	50
50	0.0	5.0	50

Table 6-2. Default Values for Histograms (Part 2 of 2)

<u>ID #</u>	<u>Low Value</u>	<u>Interval Size</u>	<u>Number of Buckets</u>
51	0.0	5.0	50
52	0.0	5.0	50
53	0	1	50
54	0	250	50
55	0.0	5.0	50
56	0.0	5.0	50
57	0.0	5.0	50
58	0.0	5.0	50

THIS PAGE LEFT INTENTIONALLY BLANK

6-6.2

CH-3

Table 6-3. Default Values for Plots

<u>ID #</u>	<u>Max Size of Plot</u>	<u>Lower Plot Limit</u>	<u>Upper Plot Limit</u>
26	Unlimited	0.	456.
27	Unlimited	0.	456.
28	Unlimited	0.	114.
59	Unlimited	0.	228.

Table 6-4. Available Report Actions and Their (Default) Values
(Part 1 of 2)

HISTG - Modify a histogram (see table 6-2)

PLOT - Modify a Plot (see table 6-3)

ON - Turn a specific report on - (all reports on except Memory Map and Out of Core Report)

OFF - Turn a specific report off - (all reports on except Memory Map and Out of Core Report)

TIME - Set a timespan(s) for reporting - (total time reported)

ALLOFF - Turn all reports off except those specified - (all reports on except Memory Map and Out of Core Report)

ALLON - Turn all reports on except those specified - (all reports on except Memory Map and Out of Core Report)

ERROR - Do not stop on an option request error - (stop on an input error)

DEBUG - Program debug requested - (no debug)

ALLOC - Stop program after a specified number of memory allocations have been requested - (entire tape processed)

NREC - Stop program after a specified number of tape records have been processed - (entire tape processed)

| NOUSER - Do not print USERID on any report - (USERID printed on certain reports)

IDLE - Turn off all Idle Monitor reports - (all IDLE reports on)

WASTED,CORE,IO,CPU,RATIO,URG - Changes parameters used in the Excessive Resource Usage Report - (20K,50K,30MIN,30MIN,5,40)

| ABORT - SNUMBs not to report in the ABORT Report - (all SNUMBs that abort are reported)

PLTINT - Change Interval at which plots are printed - (10 MIN)

FSTSLV - Change the lowest allowable user program number - (14 decimal)

MASTER - Define SNUMBs that are considered to be SYSTEM jobs - (all programs with a program number less than FSTSLV)

must be expressed as four character fields with no intervening blanks. Time is based on a 24-hour clock. If a user wants to request the time 4:07, he must input 0407. All times must include four characters.

If a start time, but no stop time, is desired, no characters should be entered after the minutes of the start time. If a stop time is requested, there must be a start time corresponding to it. If the user wants to start at the beginning of data collection and stop at some specified time, but is not sure of the start time, a start time of 0001 should be used. Figure 6-5 shows the format for this option.

6.1.11 Turn All Reports Off Except Those Specified (Action Code ALLOFF).
All reports except those explicitly identified here are to be turned off. The inputs consist of

A B C . . . Y (max of 25)

where A through Y are the report ID numbers (table 6-1) to be turned on. The format is shown in figure 6-6. This option will control the printing of all reports, including histograms if they contain a specific ID number.

6.1.12 Turn All Reports On Except Those Specified (Action Code ALLON).
All reports except those explicitly identified here are to be turned on. The input consists of

A B C . . . Y (max of 25)

A through Y are the report ID numbers (table 6-1) to be turned off. The format is the same as action code ALLOFF (see figure 6-6). This option will control the printing of all reports, including histograms if they contain a specific ID number.

6.1.13 Continue Data Reduction After an Input Option Error (Action Code ERROR). This code allows data reduction to continue when an error has been detected and reported in an input option request. The default value will abort data reduction and report the error. Only the Action Code card is required.

6.1.14 Debug For a Given Program Number (Action Code DEBUG). This is a debug option which supplies large amounts of output for a given program number. It should be used only in cases of data reduction problems. Card 1 contains the word DEBUG and card 2 contains a program number. A program number of -150 will provide detailed debug on system scheduler activities.

6.1.15 Stop After a Specified Number of Tape Records Processed (Action Code NREC). This option is useful when a tape problem occurs and the entire tape cannot be processed. When this occurs, the program will usually abort with an I/O error and some reports might be lost. If a tape error does occur during data reduction, the operator should type a "U" in

response to the operator action request made by GCOS in processing tape errors. If the operator performs this action, the data reduction program will abort gracefully.

Unfortunately, there are times when a tape error will cause a program abort without giving the operator a chance to respond with a "U". In these cases reports will be lost and this option will need to be used

THIS PAGE LEFT INTENTIONALLY BLANK

6-15.2

CH-3

Card 1 A
2 N M
3 B C D E ...

where

A = The word TIME

N = Report ID name to be time spanned (table 6-1)

M = Number of different times appearing on Card 3

B,C,D,E = Start and stop times used to define the time spans.
Times must be separated by one or more blanks.

Figure 6-5. TIME Action Code Format

6.1.21 Change the Program Number for the First Slave Job (Action Code FSTSLV). In the GCOS system, certain program numbers are assigned to system jobs. For example \$CALC is program number 1, \$PALC is program number 2, \$SYOT is program number 3, etc. In the WWMCCS system, all programs with a program number less than 14 (decimal) are considered system programs. This option allows the user to alter this program number from its default value of 14. The first card contains the word FSTSLV and the second card contains the new program number. For non-WWMCCS systems, FSTSLV should normally be set to 10.

6.1.22 Request that Certain Jobs be Considered System Jobs (Action Code MASTER). There are certain jobs executed during the course of a day which have program numbers that would designate these jobs as user jobs. However, in actuality they are system jobs and should be considered as system overhead. Examples of such jobs are VIDEO, HEALS, the GMF MONITOR, etc. This option allows the user to define up to ten jobs that should be considered as system jobs. The first card contains the Action Code MASTER. The second card contains the number of jobs to be defined as system jobs. The third card contains the SNUMB of each job to be considered as a system program. Each SNUMB must be followed by at least one blank column.

6.1.23 PALC Report Print Control (Action Code PALC). Due to the excessive amount of output possible from the PALC report, a time control can be set to print only those activities that are in any PALC state greater than the time limit. This time limit defaults to 600 seconds (10 minutes). The first card contains the word PALC and the second card contains the new time limit, in seconds.

6.1.24 Request the Special Job Memory Reports (Action Code SPECL). If the analyst desires to track the memory demands for a specified number of jobs (not to exceed ten), this input option should be invoked. This option will cause two reports to be produced. One report will indicate every time the requested job(s) was swapped or issued a MME GEMORE/GEMREL for memory, how long it was swapped, or how long the GEMORE was outstanding, and how much memory the job(s) required. In addition, every time the special job issues a MME GEMORE, a line from the Memory Map Report will be generated. This line is generated by default and is not dependent upon whether or not the Memory Map Report is enabled. When the analyst wants to match the Memory Map output to the Special Job output, he must do so based on the time value. For example, if the Special Job Report indicates that FTS issued a MME GEMORE at 16.81057, the user would then examine the Memory Map for a line of output with a time smaller than 16.81057, but where the time on the next line is greater than or equal to 16.81057. For example, the Memory Map might have a line of output with a time indication of 16.81052 where the next line of output was 16.81065. In this case, the line of output at 16.81052 shows what memory looked like at the instant in time that FTS issued the MME GEMORE. If the Special Job

Report indicates that FTS was swapped after issuing the MME GEMORE, the analyst could examine the Memory Map in order to determine why FTS was forced to swap.

A line of the Memory Map is also generated every time the GEMORE for the special job was denied or the special job was forced to swap in order for the GEMORE to be satisfied. This line of the Memory Map would show what memory looked like when the special job was denied the memory request or was swapped from memory. A final line of the Memory Map is produced whenever the special job's memory demand was met. For the swap/denied case and the memory-met case, the Special Job Report and Memory Map are matched by locating identical time values on each report. By generating the Memory Map, the analyst can determine if there are certain jobs that are preventing other jobs from acquiring required memory resources. In this case, the Special Job Report and Memory Map Report can be correlated by matching up the time values from both reports with the identical time values. This is especially useful in an analysis of the Timesharing Subsystem or the File Transfer System.

A second report will also be produced which indicates the average memory size of the job(s) during the course of its execution. This average is taken over increments of time where the time increment used, is the same increment that is used to produce the series of plots. The option consists of three cards where the first card contains the word SPECL, the second contains the number of jobs to be analyzed, and the third card contains the list of SNUMBs separated by at least one blank column.

6.1.25 Process Data on a WW6.4 System (Action Code RN). If the data reduction program is to be run on a WW6.4 system, the user must use this input option. It consists of the letters RN typed on a data card.

6.1.26 Produce a Memory Map Only Under Certain Memory Demand Conditions (Action Code MAPART). Due to the enormous amount of output generated by the Memory Map and Out of Core Reports, it is suggested that a site not produce these reports as a standard procedure. However, these reports are very useful in that they do provide a complete picture of memory as well as a total list of all jobs waiting for memory. In order to provide an analyst with the capability of obtaining these reports, without being buried in computer output, this new option has been designed. When used, this option states that a line of the Memory Map and Out of Core Reports should be generated only when the number of activities waiting for memory surpasses a certain limit. To invoke this option, a two-card format is required. Card 1 contains the word MAPART and card 2 contains the number of activities that must be waiting memory before a line of output will be generated for the Memory Map and Out of Core Reports.

Table 6-5 shows all the MUDRP file codes and their corresponding reports.

6.3 Outputs

In this section, a simple explanation of how each report was derived from the data is given. Subsection 6.1 discussed how the ranges and other options of each report may be modified to fit an individual installation. While this section will provide some insight as to how an analyst should proceed to review all the reports produced by the MUDRP, section 14 provides a step-by-step approach as to how a memory analysis might be conducted.

Immediately prior to the output of the histograms, the user will find a printout containing processing information. Included in this information is the following:

- o Printout of all input options selected by user
- o Indication of multireel tapes that are being requested and have been mounted
- o Indication of the monitors that were active during data collection
- o Error messages - all error messages are either self-explanatory or else followed by the words "For Information Only." The latter messages are used by CCTC for future enhancements and as such can be ignored by the user.
- o If the time frame option was used, and indication of when the various time frames were reached.

6.3.1 MUM Title Page. The Memory Utilization Monitor (MUM) title page contains a summary of the systems configuration and activity over the measurement period (see figure 6-9). It displays the time the monitor was initiated and terminated, as well as identifying the system which was monitored and the tape number(s) containing the data. The configuration information is augmented by the amount of memory dedicated to the operating system itself, including that used by the memory allocation program. These figures will give the user a good idea of how much hard core space remains and could be used for SSA module hard core loading. If SSA cache is also configured, the amount of memory being used for this feature is also listed. The version number should be 01-82.

Immediately following is a summary of the work processed over the measurement period. The first set of lines provides information concerning the overhead generated by the actual data collection. The monitor name is given, its CPU time in seconds, and its overhead as a function of total processor power. The GMF executive overhead is separated from the actual monitors and is listed as "EXEC". The monitor "NAME" is an area of code within the Mass Store Monitor and even though

Table 6-5. File Code for MUM Reports
(Part 1 of 2)

20	Activity Resource Report, Special Job Reports
21	IDENT Report
22	Special Job Report (temporary file)
23	Special Job Report (temporary file)
24	Urgency Over Time Report (temporary file)
26	Zero Urgency Job Report (temporary file)
27	Activity Abort Report
31	Plot 1 - (see table 6-1 for Plot Definition) (temporary file)
32	Plot 2 - (see table 6-1 for Plot Definition) (temporary file)
33	Plot 3 - (see table 6-1 for Plot Definition) (temporary file)
34	Excessive Resource Report
35	Plot 4 - (see table 6-1 for Plot Definition) (temporary file)
36	Used for outputting all plots
37	Used for outputting Out of Core Report, Memory Map, and Peripheral Allocator Report
42	Histograms, System Program Usage Report, Memory Statistics Report, Distribution of Urgency Over Time Report, Zero Urgency Job Report
45	Out of Core Report (temporary file)
51	Memory Map Report with one file required for each 128K Memory configured (temporary file)
52	Memory Map Report with one file required for each 128K Memory configured. (temporary file)
53	Memory Map Report with one file required for each 128K Memory configured (temporary file)
54	Memory Map Report with one file required for each 128K Memory configured (temporary file)

***** THE MEMORY UTILIZATION MONITOR *****

VERSION 01-82

MONITORING ON 79-05-04 STARTED AT 11:39:50 AND COMPLETED AT 13:19:56 FOR A TOTAL TIME OF 1.67 HOURS
ON SYSTEM OSCC2 RUNNING 6.4.1D OF TAPE D0002

THE SYSTEMS CONFIGURATION CO LISTED OF:

2 - 6680 CENTRAL PROCESSORS
2 - INPUT/OUTPUT MULTIPLEXORS
WITH 24 I/O CHANNELS
512 - 1024 WORD BLOCKS OF MEMORY
52 OF WHICH WERE USED BY THE HARD CORE SYSTEM ITSELF
3 OF WHICH WAS USED BY CALC
5K USED BY SSA CACHE (NOT HARD CORE)

THE SYSTEM PROCESSED THE FOLLOWING OVER THE MEASUREMENT SESSION:

MONITOR	TIME(SECS)	% OVERHEAD
EXEC	100	1.7
MUM	60	1.0
CPU	40	.7
IDLE	40	.7
TOTAL		4.1

175 ACTIVITIES WERE PROCESSED AT A RATE/HOUR OF 104.89
OF THESE 50 WERE SYSTEM SCHEDULAR ACTS
125 ACTUAL ACTIVITIES WERE PROCESSED AT A RATE/HOUR OF 74.85

128 MOVES WERE PERFORMED AT A RATE/HOUR OF 76.72
266 SWAPS WERE PERFORMED AT A RATE/HOUR OF 159.44
201466 TIMES THE PROCESSORS WENT IDLE YIELDING A
20 % IDLENESS OF THE PROCESSORS

THE FIRST PROCESSOR WENT IDLE 122479 TIMES, YIELDING 23 % IDLENESS
THE SECOND PROCESSOR WENT IDLE 78987 TIMES, YIELDING 17 % IDLENESS

THE MEMORY ALLOCATOR WAS CALLED 1907 TIMES - 0 OF WHICH RESULTED IN NO STATE CHANGE

THE TOTAL CPU TIME IN SECS WAS 6991 THE TOTAL IO TIME IN SECS WAS 15662 CPU/IO RATIO IS 0.446373
WEIGHTED MEMORY SURPLUS IN K WORDS WAS 46
WEIGHTED MEMORY SHORT-FALL IN K WORDS WAS 50 INCLUDES CALC AND PASC QUEUES

Figure 6-9. MUM Title Page Report - Idle Monitor Active

listed separately it is also included under the monitor "MSM". The Monitor "FMS" is also an area of code within the Mass Store Monitor, but in this case it has not been included under the monitor "MSM". These two special areas of code, within subroutine T7 (connect trace processing), are considered to be high usage areas and as such consume significant processing resources. In order to determine the true overhead of these areas, so that future code optimization can be considered, these areas are being reported separately.

Monitor "CM" in this report describes the processor overhead of subroutine T4 (terminate processing) and subroutine T22 (start I/O processing). Monitor "MSM" in this report describes the processor overhead of subroutine T7 (connect processing). Therefore, if the Channel Monitor was active, but the Mass Store Monitor was not, this report will still list both "CM" and "MSM" as contributing to the processor overhead. The total Channel Monitor overhead will be found by adding the overhead of the "CM" monitor, to the overhead of the "MSM" monitor, to the overhead of the "FMS" monitor.

If both the Channel Monitor and Mass Store Monitor were active, then the combined overhead of both monitors can be found as the sum of "MSM" + "CM" + "FMS".

For purposes of this report, % overhead is computed as

$$\frac{(\text{CPU TIME used by monitor})}{(\text{Total Elapsed Time}) \times (\text{number of Processors})}$$

Following this are several lines describing the work performed during the monitoring session. These lines are self-explanatory.

If a termination record is not processed, either because the monitor aborted before a termination record could be written or else time frames were used, the lines describing GMC overhead will not be printed.

The number of times a processor went idle is derived from the idle processor traces captured by the IDLEM, with the percentage of processor idle also being gathered by the collection of idle state information. This is shown system-wide (i.e., for all the central processors and then individually for each processor). This information will not be present if the IDLEM was not active or if its output reports have been disabled by a data card option (see figure 6-10).

The number of memory allocator calls, as counted by the monitor, is shown. This much less than the number of calls to the multitude of SSA modules used by the Core Allocator and consists only of those that may have altered the memory state of the system. The second figure shows how many times a memory state change might have taken place and did not. This could be caused by no allocation being possible or by a call to the allocator pertaining to a matter other than allocation (i.e., a console message).

6.3.3.20 Report 20 - The Elapsed Duration of User Activity in 10ths of a Second. This report presents the clock time that the allocator knew of a user activity's existence, measured from its first memory demand to its termination. This includes all time spent in a GEWAKE, in memory, and swapped.

For this report, an entry is made for each user activity that terminates. See report 10 for an explanation of user vs. system activities.

6.3.3.21 Report 21 - The Total Elapsed Time a User Activity Was in Memory. This report shows the duration of elapsed clock time each user activity had memory allocated to it. It helps describe the system workload requirements.

For this report, an entry is made for each user activity that terminates. See report 10 for an explanation of user vs. system activities.

6.3.3.22 Report 22 - The GEMORE Service or Denial Time - 1/10 Second, Elapsed. The time from a GEMORE request until the activity is allocated the extra memory, swapped to achieve the additional memory, or denied the memory is displayed in this report.

For this report, an entry is made for each activity whose GEMORE request is not longer present.

6.3.3.23 Report 23 - The Request Size of GEMOREs. All GEMORE requests are shown in this report with the displayed size in 1K blocks.

For this report, an entry is made for each GEMORE request.

6.3.3.24 Report 24 - Delay Time in the System Scheduler. The amount of time a job spends in one of the scheduler queues is displayed in this report. The Allocation Status Report can be used to display particular jobs that are delayed for excessive amounts of time.

6.3.3.25 Report 25 - Delay Time Until Core Allocation. This report displays the total amount of time activities spent in the various allocation phases prior to core allocation. The Allocation Status Report can be used to display particular activities that are delayed for excessive periods of time.

6.3.3.26 Reports 26 through 31 - The Elapsed Time of a Busy State of the Processors. These reports present the elapsed clock time between the idle states of each individual processor. The reports supply an indication of how each processor is utilized versus the others in the system.

For these reports, an entry is made at each idle state of a processor. IDLEM data is used to produce these reports. These reports will not be produced if the IDLEM was not active or if the IDLEM reports have been disabled by user input option.

6.3.3.27 Report 32 - The Elapsed Time of a Busy State of Processors. The elapsed clock time between idle states of all processors is presented in this report.

For this report, an entry is made for each processor idle state. IDLEM data is used to produce this report.

6.3.3.28 Report 33 - Elapsed Time Between Allocator Calls in 1/100 of a Second. This report shows the elapsed clock time between calls to the allocator and shows if the allocator is receiving sufficient service.

For this report, an entry is made for each user activity that terminates.

6.3.3.29 Report 34 - The I/O Time Charged per User Activity in Seconds. This report indicates the I/O time charged to each user activity.

For this report, an entry is made for each user activity that terminates.

6.3.3.30 Report 35 - The CP Time Charged per User Activity in Seconds. This report presents the CP time charged to each user activity. For this report, an entry is made for each user activity that terminates.

Reports 34 and 35 report the total CPU and I/O times used by a user activity while the monitor was active. These histograms are not generated for programs with program numbers less than 14 (i.e., system programs). See report 10 for additional user options in defining system activities and user activities.

6.3.3.31 Report 36 - The Number of Times a User Activity was Swapped. This report shows the swap count per user activity. The total number of swaps a user activity incurs is the user argument, as counted by the monitor. See report 10 for additional user options in defining system activities and user activities.

For this report, an entry is made for each user activity that terminates.

THIS FIGURE HAS BEEN DELETED

Figure 6-14.

6.3.3.32 Report 37 - The Total Elapsed Time a User Activity Was Swapped. This report indicates the total time a user activity was inactive due to a swap. After each swap is completed, an accumulator is updated, and if an activity is terminated, an entry is made to this report. See Report 10 for additional user options in defining system activities and user activities.

6.3.3.33 Report 38 - The Number of Times a System Activity Was Swapped. This report is the same as Report 34 except for system activities. See Report 10 for additional user options in defining system activities and user activities.

6.3.3.34 Report 39 - The Total Elapsed Time a System Activity Was Swapped. This report is the same as Report 35 except for system activities. See Report 10 for additional user options in defining system activities and user activities.

6.3.3.35 Report 40 - Number of Extra Activities That Might Fit in Memory Using Compaction. This report shows how memory might have been used more optimally. It takes the total amount of available memory (displayed in Report 5) and attempts to fit in those activities waiting memory. If an activity fits, the memory available is decreased, and the next activity is tried. If an activity does not fully fit, the next activity is tried. This continues until all available memory is used or until all the activities waiting have been tried. The search starts at the first waiting program and progresses serially down the program numbers of those waiting. This search ignores the actual size of "holes" or quadrant-crossing and is not necessarily obtainable or optimal. For this report, an entry is made at each allocator call.

6.3.3.36 Report 41 - Number of Extra Activities That Might Fit Memory Without Compaction. This report is the same type as Report 40. In this case, activities are fit into existing holes and are ordered by urgency. The search progresses down the activities serially, beginning at the highest urgency activity. This histogram presents a good picture of how well the core allocator is performing its function.

For this report, an entry is made at each allocator call.

6.3.3.37 Report 42 - The Percent of Size-Time Product Used by a User Activity. This report shows the percentage of user each activities' size-time product over its run-time duration. An entry is made for each user activity that terminates. See Report 10 for an explanation of user and system activities.

6.3.3.38 Report 43 through 49 - The Length of Idle State in the Processors. The elapsed clock time of an idle state is given in these reports for each individual processor and also as an average for all processors. They supply an indication of how each processor was utilized versus the others in the system. They also provide information on how busy the processors are. These reports should be used in conjunction with Reports 26 through 32.

statement that the \$LIMITS card appears to be requesting more memory than is actually required by this job. The user should be questioned in order to determine if this is in fact true. In the Honeywell System, a user will receive whatever amount of memory requested on the \$LIMITS card, whether or not the amount of memory is actually needed. The Ratio column shows the ratio of the total elapsed time for an activity divided by the total memory time for the activity. This value gives an indication of the activity lengthening factor; i.e., how run time is affected by resource contention. For those activities using excessive memory, the report also indicates, under the MEM MIN column, the amount of time the activity was in memory. The value being used for the urgency check is the average urgency recorded for the activity and not the maximum urgency of the activity. The default values for an entry being made to this report are listed in table 6-4. These values can be changed via a previously described input option. This report will be produced whenever the Activity Resource Report is produced and will be turned off whenever the Activity Resource Report is off (see figure 6-21).

6.3.11 Allocation Status Report. This report will track an activity as it proceeds through different phases of Allocation. The report will list the SNUMB-Activity #, amount of memory the activity will require, its current status, the time it entered that phase of allocation, the time is completed that phase of allocation, the total time spend in a given phase of allocation, the device type it is waiting for, and the number of devices the activity is waiting for. Due to the manner in which data is collected for this report, it is possible that certain phases of allocation will be missed, especially if that phase of allocation occurs within a short time span. This report will give a good indication of how long it is taking activities to pass through the various allocation phases prior to core allocation. Following is a list of the more common phases of allocation and their meanings:

- New Act - Activity has just entered the Peripheral Allocator
- Wait Media - Activity is waiting for a device
- Wait Mnt - Activity is waiting for a patch or tape to be mounted
- Core Queue Full - Activity has been completely processed and is waiting for the Peripheral Allocator to send the job to the core allocator
- Alloc Done - Activity has been sent to core allocator. For this case the stop time and total time columns have no real meaning. These columns simple are reporting the amount of time it took the monitor to realize that the activity had reached the core allocator
- LIMBO - Activity is in Limbo and has not even been granted permission to run
- HOLD - Activity is in Hold and has not even been given permission to run
- SCHED - Activity was in one of the System Scheduler queues.

Only activities found to be in a state for more than 600 seconds will be reported. This limit can be changed by using the PALC input option. See figure 6-22 for a sample of this report.

6.3.12 Plot Reports. Three different plot reports are produced by the data reduction program. All plots are produced under 10-minute intervals, where the interval can be modified by the user. At every allocator call, the various parameters to the plots are accumulated and every 10 minutes, the accumulated parameters are averaged and an average value is output

THIS PAGE LEFT INTENTIONALLY BLANK

6-53.2

CH-2

EXCESSIVE RESOURCE USAGE REPORT ON SYSTEM DNAH66 ON 81-12-17 AT TIME 10:01

SNUMB-ACT	WASTED MEMORY	MEM USED	IO SECS	CPU SECS	RATIO	MEM MIN	MAX URG	USERID	IDENT
2027Q- 1		56				1.2		FCCCSE	1020, WA024, CSE, 250, MCCANN, SECRET
2044T- 1		81				2.8		FCCCSE	1010, WA803, CSE, 264, UARRETT, UNCL
2057T- 1		64				3.4	47	FCCCSD	1010, NS0033, CSD, 283, PADILLA, UNCL
2083Q- 1		56				1.0		FCCCSE	1020, WA024, CSE, 250, MCCANN, SECRET
2085T- 1		92				5.0	55	DBA	1020, TRAX.A, CSA, 254, MATHEWS, SECRET
2106T- 1		81				2.7		FCCCSE	1010, WA803, CSE, 264, BARRETT, UNCL
SR01 - 5		65				0.6		FCCCS	1030, SR01, CSC, 250, MCCANN, TOPSEC

Figure 6-21. Excessive Resource Usage Report

ALLOCATION STATUS REPORT						
SNUMB	MEMORY	STATUS	TIME IN	TIME OUT	TOTAL TIME(SEC)	NUMBER
7338T- 1	10	WAIT MNT	12:42	12:52	600.	
7338T- 4	51	CORE QUE FULL	12:46	12:57	660.	
29675- 2	36	WAIT MEDIA	12:49	13:00	660.	
					DSS191	2

Figure 6-22. Allocation Status Report

to the plot. Each horizontal position has a delta value, which is printed on the plot. The delta value is computed from the following formula:

$$\frac{(\text{Upper Plot Limit} - \text{Lower Plot Limit})}{114}$$

114

The plot limits can be set by the user, with the default values shown in table 6-3. If the user changes the maximum limits, the new maximum limit selected should be divisible by 114. If a plotted variable is beyond or on an axis limit, it will be positioned at the axis limit. If any 2 points coincide, the position of coincidence will be marked with a 2. If 3 points coincide, the position of coincidence will be marked with a 3. (See figure 6-23). The end of the plot will contain a summary of the minimum and maximum values of each curve.

In figure 6-23, we see that there was no memory shortfall at all between 12:49 and 14:09 (points A & B are coinciding with the left axis; i.e., a 2 is output). At 14:19 and 14:29, both curves continue to coincide, but there is now a shortfall of 48K (12th point times delta of 4). At 14:39, the memory shortfall of the CALC increased to 92K but the memory shortfall of the CALC plus the PALC queue increased to 116K.

In order to obtain a continuous curve from the plots, the user needs only to connect the corresponding letter points (see figure 6-23).

6.3.12.1 Plot 1 - Available Memory vs. Outstanding Demand in Core Allocator Queue vs. Outstanding Demand in Core Allocator Queue + Peripheral Allocator Queue. This three parameter plot provides an overview of the time dependence of both the system load and memory availability. It can aid in better balancing the workload across the day and in determining when memory shortfalls or surpluses exist. The addition of memory demand waiting in the Peripheral Allocator is an attempt to give a truer picture of how much additional memory could be properly utilized, if available. As long as the B and C points fall to the left of the A points, a memory surplus exists. If the B and C points fall to the right of the A points, a memory shortfall is present.

6.3.12.2 Plot 2 - Memory Shortfall in Core Allocator vs. Memory Shortfall in Core Allocator + Peripheral Allocator. This plot is obtained from the previous plot by simply calculating the actual shortfall and plotting the shortfall points.

6.3.12.3 Plot 3 - Number of Activities in Core Queue vs. Number of Activities in Peripheral Allocator Queue. In this plot, the number of activities waiting for memory is displayed, instead of their memory demand.

6.3.12.4 Plot 4 - Average Size of TSS, FTS, NCP. This plot displays the average size of TSS, FTS and NCP as they change their demand for memory over time. These three programs are those whose memory demand would significantly change over time. The FTS and NCP programs are part of the WIN subsystem.

MEMORY STATISTICS TABLE (DIAGNOSIS)

DATE	START	STOP	DURS	CPU/IO RATIO	AVG ACT SIZE(K)	USER MEM AVAILABLE	SYSTEM MEMORY	EXCESS MEMORY (+/-) CALC PALL	TIME USER ACT SWAPPER
011217	1001	1130	1.48	0.336	23.7L	21d	294	67 86	0.010 SEC
AVG # USER									
ACT WAITING MEM									
1.000									
AVG # SYSTEM									
ACT WAITING MEM									
0.									
AVG # USER									
ACT IN MEM									
3.765									
AVG # SYSTEM									
ACT IN MEM									
8.771									
RATIO OF DURATION									
VS MEMORY TIME									
1.150									
# SLAVE									
MEM USED									
75.757									
USLM ACT/HOUR									
(THROUGHPUT)									
121.770									
ACT/HOUR									
(THROUGHPUT)									
177.920									
# USER									
SHAPS									
125									
DURATION OF									
USER ACTIVITY									
70.050 SEC									
WAIT TIME FOR									
ORIGINAL ALLOCATION									
25.500									
SWAPS PER									
HOUR									
185.361									

Figure 6-24. Memory Statistics Report

SPECIAL JOB MEMORY DEMAND REPORT ON SYSTEM NMCC ON 82-06-18

TIME	DEMAND TYPE	TIME TO SATISFY IN .1 SEC	SNUMB	SIZE OF DEMAND IN K	CURRENT SIZE IN K	CALL
11.78394(114703)	GENORE	-99999	FTS	1	35	82
11.78417(114703)	MET	0	FTS	1	36	83
11.78578(114709)	GENORE	-99999	FTS	4	36	86
11.78628(114710)	DENIED	10	FTS	4	36	88
11.78633(114710)	SWAP	2	FTS	40	40	91
11.79422(114739)	GENREL	-99999	FTS	-1	39	93

**** SINCE LAST PRINT OUT FTS HAS ACCUMULATED 205 .1 SEC WAIT TIME, CURRENT TIME IS 1148

Figure 6-25. Special Job Memory Demand Report

[illegible]

Figure 9-11. Session Length Report

experiences a response time greater than or equal to the requested limit. The information printed includes Terminal ID, subsystem name, response time in seconds, and time of day. Refer to figure 9-12.

9.5.7 User Think Time Limit Report. This report is produced only if the user requests it with the THINK input option (subsection 9.6.7). This report will print out a line of information every time a terminal experiences a response time greater than or equal to the requested limit. The information printed includes Terminal ID, subsystem name, think time in seconds, and time of day. Refer to figure 9-12.

9.5.8 Terminal Session and High Terminal Usage Reports. The Terminal Session Report (figure 9-13) is produced whenever the Statistical Summary Reports are requested. The report gives an account of every session that occurs during the monitoring session. Every time a user logs off or is logged off due to a DN355 abort, TCALL, or monitor termination, an entry into this report is produced. The report gives the Log On Time, Log Off Time, Terminal ID, USERID, Subsystem, Subsystem Size (MIN MAX), Session Length, Response Time, # Inputs, # Outputs. If a terminal was logged onto a subsystem when CAM started, then there is no immediate way for CAM to determine the subsystem being used by the terminal. In this case, the CAM data record will be checked to see if a USERID exists for the terminal. If one does exist, it means the terminal is logged on to TSS and the log on name is changed from UNKNWN to TSS. If there is no USERID for this terminal in the record, the subsystem name is set to UNKNWN. This occurs for the WIN lines and any user logged on to VIDEO when the monitor is started. If a user JDACS to a new subsystem, CAMDRP will disconnect the current line, calculate all statistics and reconnect the line to the new subsystem. Both the USERID and Subsystem Sizes are given only for TSS users. It has been noticed that some WIN pseudo terminals appear logged onto TSS, but have no USERID and no input/outputs. The cause for this discrepancy is under investigation. The Session Length is given in seconds. The Response Time is given in seconds and is the average response time over the session. The # Inputs is the number of input requests sent by the user. The # Outputs is the number of output response groups sent to the user. This report can help pinpoint excessive response times. It can also be used to determine if a terminal is logged onto the system and is not being used (low inputs, high or low outputs, long session length).

The High Terminal Usage Report is included as part of the Terminal Session Report and provides a list of terminals that have been logged on for a specified percent (default 75%) of the session. This will list terminals by ID and type, give the percent of time the terminal was logged on, the number of sessions during this time, the number of inputs and the number of outputs from the terminal. (See figure 9-14.)

9.5.9 Opcode Count Report. This report (figure 9-15) is produced whenever the System Summary Reports are produced. This report gives a listing of all the opcodes that were transmitted between the H6000 and the DN355, and a count of how many of each opcode there were. This report is of interest mainly when the following opcodes appear:

EXCESS THINK/RESPONSE TIME REPORT FOR NMCC2 ON 122181

TERMINAL C6 ON	SUBSYSTEM TSS	HAD A THINK TIME OF 208 SECS AT 16:35:55
TERMINAL C1 ON	SUBSYSTEM TSS	HAD A THINK TIME OF 120 SECS AT 16:48:21
TERMINAL BT ON	SUBSYSTEM TSS	HAD A THINK TIME OF 223 SECS AT 16:48:57
TERMINAL BD ON	SUBSYSTEM TSS	HAD A THINK TIME OF 197 SECS AT 16:49:11
TERMINAL BT ON	SUBSYSTEM TSS	HAD A RESPONSE TIME OF 50 SECS AT 16:43:54

Figure 9-12. Response Time/User Think Time Limit Report

TERMINAL SESSION REPORT FOR NMCC2 ON 102182 AT 14:18: 0.000

LOG ON TIME	LOG OFF TIME	TERMINAL ID	USERID	SUBSYSTEM	SUBSYSTEM MIN	SUBSYSTEM SIZE MAX	SESSION LENGTH	RESPONSE TIME	NUMBER INPUTS	NUMBER OUTPUTS
18:53:17	18:54:37	CN	DJ3JI32411	TSL	1K	3K	80 SEC	4.500 SEC	4	8
15:01:25	18:54:37	4C		SEND			13993 SEC	0. SEC	0	0
18:48:31	18:54:37	AA	DJ8XC344JG	TSL	1K	26K	367 SEC	3.324 SEC	39	125
18:03:07	18:54:37	CV	DJ3N13203	TSS	OK	29K	3090 SEC	10.120 SEC	103	499
14:18:58	18:54:37	AB		UNKMWN			16539 SEC	0. SEC	0	4618
18:42:07	18:54:37	AP	DJ3JI32401	TELNET			751 SEC	8.844 SEC	32	69

Figure 9-13. Terminal Session Report

second, or 1.76 seconds. The shape of the distribution is symmetric -- about the same number of activities had values over 1.76 seconds as under 1.76 seconds, and the standard deviation is small when compared to the average. The absence of a second line under the "Entries Total" line indicates that no activities stayed in memory longer than 2.4 seconds. When a distribution resembles this example, use the "Average" printed at the bottom as the Representative Value.

14.6.2.2 Skewed Distribution. Figure 14-3 shows another distribution. This distribution is "skewed" (i.e., not symmetric), because most of the activities spent around 0.1 to 0.7 seconds in memory, while some spent as much as four or five seconds. Care should be taken when selecting a "Representative Value" from this distribution. If the analyst wants to emphasize the "typical" activity, which stayed in memory 0.3 seconds or less, he could select the "median" (the value which evenly divides the activities in the distribution -- half spent less time in memory, and half spent greater time in memory). In figure 14-3, the median is about 0.29 seconds. The median can be estimated from these reports by descending down the "Cumulative" column (not displayed in the figure) until the value first exceeds 50. The median falls within the time range of this row.

14.6.2.3 Distribution With Outliers. There are some instances when the distribution will also have some values that were too big to fit in the histogram. This condition will be indicated by an additional output line at the bottom of the report. This line will indicate the number of occurrences that were outliers, the average for just the outliers, and the average for the values that fit into the report, minus the outliers. The three important factors about these types of distributions are: (1) the amount of times they occur; (2) the percent of the total values that are outliers; (3) the "in-range average."

If the percent of outliers is greater than 10%, the analyst should use the overall average, given in the first line of the report, for any comparisons. If the percentage of out-of-range values is less than 10%, the analyst can use the "in-range average" value for his comparisons since the effect of the outliers will most likely be minimal on the total system performance.

14.6.3 Memory Evaluation. The first step in a memory evaluation is to summarize all the pertinent information. The easiest way to do this is for the user to create a table similar to figure 14-4. The information for each column is collected from various MUM reports. Once the chart is filled out, the analyst can then ascertain a reasonable idea of the overall memory status of the system. The reports required to collect the statistics will be discussed, as well as an analysis procedure. The current version of the MUM will automatically produce this table (see subsection 6.3.13 - Memory Statistics Report).

DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 14:14:24 ON 81-05-15

The Total Elapsed Time An Activity Was In Memory

		Percent of Occurrence											
Tenth	Second	00	05	10	15	20	25	30	35	40	45	50	
		/	/	/	/	/	/	/	/	/	/	/	
0-1		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx											
2-3		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx											
4-5		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx											
6-7		Ixxxxxxxxxxxxxxxxxxxx											
8-9		Ixxxxxxxx											
10-11		Ixxxx											
12-13		Ixxx											
14-15		Ixx											
16-17		Ix											
18-19		Ix											
20-21		I											
.													
.													
56-57		I											
118669 Entries Total		Average 4.247 Variance 18.52 Standard Deviation 4.3											

Figure 14-3. Sample Skewed Distribution

DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 14:14:24 ON 81-05-15

The Total Elapsed Time An Activity Was In Memory

		Percent of Occurence										
Tenth	Second	00	05	10	15	20	25	30	35	40	45	50
		/	/	/	/	/	/	/	/	/	/	/
0-1		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
2-3		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
4-5		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
6-7		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
8-9		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
10-11		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
12-13		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
14-15		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
16-17		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
18-19		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
20-21		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
.		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
.		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
.		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										
56-57		Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx										

118669 Entries Total Average 4.247 Variance 18.52 Standard Deviation 4.3

Figure 14-3. Sample Skewed Distribution

If Column 7 shows a surplus of memory, but does not exceed the aforementioned limits, the implication is that the current system has sufficient memory, but that the system is approaching memory saturation. If Column 7 shows a shortfall of memory, and the value is greater than 15% of the total available memory or greater than two times the value reported in column 4, then the implication is that memory is a constraint on this system. Finally, if Column 7 shows a shortfall of memory, but does not exceed the aforementioned limits, the implication is that the system has reached memory saturation, but is still able to process the current workload.

Step 2 - It should be stressed that the value reported in column 7 is calculated over the entire measurement period and therefore could be biased by periods of heavy or light activity. It is for this reason that the user is urged to run the monitor during those periods of time where the workload is considered to be heavy and constant. In order to determine if the above type of biasing is occurring, the user may want to check Plots 1-3. If it appears that there is a mixing of light processing and heavy processing the user may want to re-run the data reduction program, using the time-frame option, to separate the heavy processing time.

Step 3 - Calculate the ratio of column 5 divided by column 4. This is an indication of the maximum number of user jobs that your system can support at any one time, without the occurrence of significant swapping. If the value in column 10 is equal to or exceeds this ratio, then the implication is that the system has reached memory saturation. If the value in Column 10 is within 2 units of the ratio, then the system probably has sufficient memory but is approaching saturation. Finally, if the value in column 10 is less than the ratio by more than 2 units, the current system has sufficient memory.

This step can be further verified by checking columns 14, 16 and 18 for indication of significant swapping.

Step 4 - If column 13 is less than 85, the current system should have sufficient memory and the other steps should not be showing indications of memory problems. If column 13 is between 85 and 95 then the current system is approaching saturation and at times may be showing some indications of a backlog. If the figure exceeds 95, then other steps should be indicating signs of moderate to severe memory problems.

Step 5 - If column 8 is greater than or equal to 3, the indication is that memory has become a constraining factor.

Step 6 - If Column 12 is greater than 2, the indication is that memory wait time is high and that memory is probably a constraining factor.

At this point, the user should have a fairly good indication as to whether or not memory is a constraining factor. The following steps will indicate some additional reports that the user should reference to determine those jobs that might be causing the memory problem.

Step 7 - One of the largest users of resources are jobs that abort and then must be rerun. Aborts usually occur due to user errors, but hardware aborts are not uncommon. If management is aware of aborting jobs and the reasons for them, they can possibly save substantial system resources. The Abort Report is described in subsection 6.3.8 and gives an indication of the system resources being wasted by aborting jobs.

Step 8 - It is important for management to be aware of jobs that are either misusing system resources or are requesting large amounts of system resources. Upon identifying such jobs, these jobs could be redesigned, scheduled for non-peak processing, or, in the case of wasted resources, the waste could be eliminated. The Excessive Resource Report allows the user to uncover jobs of this type and is described in subsection 6.3.10. When using this report, the following are suggested parameter values:

wasted core - 5K
memory - either 35K (WWMCCS standard) or 2 times the value in column 4
CPU time - 15 minutes
IO time - 30 minutes
URG - 40
RATIO - 2

Step 9 - By examining the System Program Usage of Memory Report, the user can determine those system type jobs that are requiring memory. It is possible that some of the system jobs can be eliminated or at least reduced in size. This is especially true for the TSS. However, it must be realized that a limitation on Time Sharing size may adversely effect TSS response. In many cases, if large file transfers are being processed during prime time, the size of the FTS WIN subsystem can rise to 70 or 80K. By not allowing WIN file transfers to run during prime time, significant memory savings can result.

Step 10 - As is explained in great detail in subsection 6.3.15, it is vitally important that the overall urgency level of jobs being processed remain low. The Distribution of Urgency Report can be used to determine the overall urgency level of jobs being processed. This report should show that 60 percent of the jobs being processed at any one time have an urgency level below 40 and that a substantial proportion of these should have an urgency level between 5-10. The summary at the bottom should indicate that 75-80 percent of all activities processed had an urgency level below 20.

If this report indicates a large percentage of high-urgency jobs, then the SNUMB/IDENT report, or the Excessive Resource Report, can be used to identify those particular activities processing with a high urgency.

Step 11 - If the analyst wants to track the memory performance of a given set of jobs, the use of the SPECL input option and the generation of the Special Job Memory Reports will prove sufficient for detailed memory tracking. This procedure is especially useful in analyzing the memory requirements of TSI, FTS and the special JDA-developed software (JDSIP, JDSUP). Refer to subsections 6.1.24 and 6.3.14 for complete descriptions of these Special Job Memory Reports.

Step 12 - Another indication of poor system performance possibly caused by memory shortfall, tape drive shortfall, poor operator performance or a poor system scheduler design is the long delay of jobs as they pass through the various allocation phases prior to core allocation. The Allocation Status Report, the System Scheduler Delay Time Histogram and the Delay Time Until Core Allocation Histogram can all be used to determine which jobs, and how many jobs, are being significantly delayed during the various allocation phases. These reports are all fully discussed in section 6.

Step 13 - Memory problems may also be occurring as a result of jobs being delayed due to CPU constraints or I/O constraints. In these cases, jobs tend to sit in memory due to a lack of other system resources. Because these jobs are being delayed, other jobs cannot enter memory, and memory demands begin to backlog. Therefore, if memory is a constraint, the user should consider conducting a CPU analysis as well as an I/O analysis.

14.6.4 CPU Evaluation. The CPU evaluation will determine the general utilization level of the processor and then determine if the CPU is dominated by GCOS or user program execution. A CPU data reduction is required for this evaluation. It is also beneficial to have an associated MUM data reduction available for the same time period. Figures 14-5 and 14-6 are sample table formats that may be used to display the gathered data.

14.6.4.1 Data Recording. For Figure 14-5, data for columns 1 and 14 can be obtained from the heading page. Data for columns 2-8 can be obtained from the last section of the CPU Time Report. This report is produced every 10 minutes of elapsed time and the data of interest should be found in the last 10-minute report. Columns 9-13 are filled from Histogram Reports 1, 2, 5, 6, 9 respectively. For figure 14-6, the data may be obtained directly from the last line of the CPU Plot.

THIS PAGE LEFT INTENTIONALLY BLANK

14-18.2

CH-3

#1	#2	#3	#4
Date-Time	% Less Than 25	% Between 25-50	% Greater Than 50

Figure 14-6. CPU Processor Availability

Step 7 - The CPU Time Report and WIN Report can be used to determine those periods of time when TSS and/or WIN programs are using excessive amounts of processor time or, on the other hand, do not appear to be requesting sufficient CPU service.

14.6.5 I/O Evaluation. The I/O evaluation will determine whether the mass storage subsystem, or tape channel subsystem is the cause of system degradation. This evaluation requires the user to have processed the Mass Storage Monitor and Channel Monitor data reduction programs.

14.6.5.1 Data Recording. All output from the Mass Store Monitor and Channel Monitor are required. No individual work tables are required, but the user may generate some if he feels that it will help in his analysis.

14.6.5.2 Evaluating the Data. Chapters 7 and 8 provide a fairly detailed description of the procedure to be followed in analyzing the associated reports. In this section, reference will be made to those chapters indicating actual data values that should be used as a reference for comparison.

Step 1 - Read subsection 7.3 and subsections 8.2 and 8.3.

Step 2 - Check the crossbar configuration using the procedure described in subsection 8.2.

Step 3 - Examine the Proportionate Device Utilization Report produced by either the MSM or CM. Check for devices which have significantly higher utilization than other devices in the system. These devices are potential bottlenecks and should be more closely analyzed. It is desirable, even though perhaps not possible, to have equal utilization across all disk packs. Read subsections 7.5.18 and 7.5.20 for further details on this step. Once a pack(s) is identified, further analysis should be performed to determine the actual files being referenced on the pack (see subsection 7.6.1).

Step 4 - The histogram displaying Data Transfer Sizes for TSS Swap Files can give a strong indication of the sizes of TSS subsystems being used. TSS subsystems of over 25K can cause significant increase in overall TSS response, especially if several of these subsystems are being executed simultaneously. If more than 20 percent of the entries in this report fall in the bucket ranges above 25000, this is a strong indication that TSS response might be a problem. This problem can be further confirmed with the CAM.

Step 5 - Seek Elongation - Subsections 7.5.6 and 7.5.7 describe in detail the reports used to investigate seek elongation problems. An average seek of over 50 cylinders for DSS19ls and 100 cylinders for DSU450s should be considered significant.

Step 6 - Analyze the Channel Monitor Idle Report. This report can be generated only if the Idle Monitor was run in conjunction with the Channel Monitor. If the "% of Idle Time During Which I/O Was Active" value exceeds 25%, then substantial benefit may be obtained by eliminating I/O contention. The above value is an indication that even though the CPU is going idle (i.e., has no useful work to perform) there really is potential CPU work available. However, under current conditions, this potential CPU work is being delayed because of I/O contention.

Even though the above figure exceeds 25%, the system may not have sufficient CPU power available to handle the increased work generated by removing the I/O contention. Therefore, the analyst should also check that the "Average System % Idle" figure exceeds 15%. If this proves to be the case, then removal of any I/O contention should prove beneficial. On the other hand, if the figure is lower than 15%, then removal of any I/O contention will probably result in additional CPU contention. The Idle Report will also indicate those devices causing most of the contention. Make a record of the device numbers.

Step 7 - Examine I/O queue length and I/O queue time histograms for individual devices and channels. Queues greater than one and queue times greater than 15 MS should be considered significant. Record those devices with high contention.

Step 8 - If certain devices have been determined as bottlenecks under the procedures described in Steps 1 and 2, the Job Conflict Report should be obtained for those devices following the procedures described in Chapter 8.

Step 9 - Execute the Mass Store Monitor Data Reduction Program. Following the procedure described in subsection 7.6.1 for monitoring a specific device, the analyst should be able to determine the exact files that are causing the contention found under the earlier steps.

Step 10 - Using the CM, it is possible to perform a detailed analysis on channel queuing for a particular job. Details for this procedure can be found in subsections 8.5.13, 8.5.14 and 8.6.10.

Step 11 - This step outlines procedures for relocating files identified as candidates for file relocation. Because of automatic load-leveling activity by the GCOS operating system, an analyst has only limited flexibility for the placement of system, permanent, and temporary files:

- a. System Files. The device name on which a system file is to be placed can be specified at system startup. Care should be taken to insure that multiple high-used system files are not placed on the same disk device. If possible, separate high-use system files across disk subsystems. In addition, ensure that SSA Cache memory and FMS cache are enabled to reduce disk I/O activity to certain system files. Details for this analysis can be found in subsections 7.5.9, 7.5.10 and 7.5.19.

AD - A 124 071

GENERALIZED MONITORING FACILITY USERS MANUAL(U) COMMAND
AND CONTROL TECHNICAL CENTER WASHINGTON DC
B WALLACK ET AL. 01 DEC 82 CCTC-CSM-UM-246-82-CHG-3


2/2

UNCLASSIFIED

F/G 9/2

NL

END
DATE
FILMED
2-6-8
DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

THIS PAGE LEFT INTENTIONALLY BLANK

14-22.2

CH-3

b. Permanent Files. The device name for a permanent file can be specified at creation, whether through FMS or the ACCESS subsystem of Time Sharing. Files can be moved by changing their names, creating a new file with the old name, and moving the data. The new file can be created with a DEVICE specification.

c. Temporary Files. The device name for a temporary file can be specified in the second field of the \$FILE card in the job control deck. Jobs which run frequently can have their \$FILE cards changed. Other jobs can be controlled by policies governing the use of \$FILE cards.

Additionally, sites that have different device types may specify preferred device types to be used for temporary files. This procedure will allow activities requiring disk storage to take advantage of higher speed devices.

| Step 12 - This step identifies possible seek contention problems attributable to inadequate temporary file space. This procedure uses the SPUTIL feature of the GCOS FMS facility and the Disk Fragmentation Report available at most sites. If such a report is not available, contact CCTC/C751. It is necessary to analyze temporary disk capacity on all disk units rather than just the units identified in previous tuning steps. This analysis is necessary because the disk units exhibiting high activity due to temporary file use often have more available temporary space. The increased utilization of these disk units may be caused by inadequate temporary storage on other disk devices. For this analysis a form as shown in figure 14-7 may prove useful.

a. Report Values. The SPUTIL Report contains the following information on each disk device: (1) device identification, (2) overall capacity, (3) available disk space, (4) disk space dedicated to permanent storage. The Disk Fragmentation Report contains additional information on each disk device: (1) number of disk fragments, (2) average fragment size, (3) maximum fragment size, (4) percentage distribution of fragments by size, and (5) total fragmented space.

b. Form Entry. Enter the device identification for each disk device on the Temporary Storage Test Form. For each disk device enter: (1) the LLINK capacity as indicated by the SPUTIL Report in the Total Capacity column; (2) the temporary capacity from the SPUTIL Report in the Temp Capacity column; (3) the number of fragments from the Disk Fragmentation

SECTION 15. TIMESHARING EMULATION AND RESPONSE TIME SYSTEM (TEARS)

15.1 Introduction

The Timesharing (TSS) Emulation and Response Time System (TEARS) is a collection of FORTRAN programs which sequentially processes data the Timesharing Subsystem Monitor (TSSM) collected and written to tape. There are three separate and independent phases of reduction: response time analysis, TSS emulation, and optionally, a formatted dump. TEARS produces a number of reports showing the user load supported by TSS, computer resources expended, user response times, and service interruptions. A list of reports and report descriptions is presented in subsection 15.5.

The TSS Monitor and its companion data reduction package, TEARS, provide information to:

- o Identify those time periods in which TSS response to the user was abnormally poor.
- o Identify which resources -- CPU, I/O, memory, or non-TSS competition -- contributed to the poor response.
- o Localize the agent of poor response down to the users command, user or subsystem.

TEARS uses two inputs. The first is the data tape produced by the TSSM in the General Monitor Collector (GMC). The second input is a set of report option control cards used to alter the reports in a way other than the standard default. The various options and their formats are described in subsection 15.6.

Caveat: Section 15 is a general description of TEARS, its reports, and run-time options. It is not intended as a users manual for analysis. Little is provided here as "rules-of-thumb" toward identifying a TSS session showing abnormal behavior or as guides to interpreting the reports. An extension to section 15, in preparation but not currently available, will explain how to interpret the reports and perform an analysis.

15.2 Data Collection Methodology

The TSSM in the GMC uses trace type 74 to collect data about the TSS in execution. TSS transactions are trapped at 104 points within TSS and at one point external to the TSS; each trap produces a trace 74. Section 5 provides a description of the trace mechanism and of the 105 trace subtypes. With the information thus collected, TEARS maintains a list of active users and correlates TSS transactions with a particular user. The interaction between a user terminal and TSS is traced as is the interaction between GCOS and TSS. Processor allocation, memory allocation and user disk I/O are also sources of data.

15.3 Analytical Methodology

TEARS is a multipass program that requires analysis after each pass. The first pass displays user response time (technically, duration of line idle -- from the H6000 side of the DATANET). That display is a collection of charts, each of which shows a demand or response characteristic from minute to minute over the duration of the run. Using the charts* an analyst can determine selected timeframes of poor response for a second pass analysis of the data tape.

The second pass analysis mimics or emulates the TSS process in a limited fashion -- it remembers, dynamically, what each TSS user was doing or attempting to do, how long the attempt was ongoing, the obstacles encountered in the attempt, and some of the computer resources and TSS executive services used in the attempt. Additionally, the second pass identifies events that were of severe import or were catastrophic to the user. Several charts, histograms, and tabular reports provide both summary and snapshot of the mimic phase. Sufficient data is presented for a manual or deductive completion of the analysis.

TEARS will generate additional reports if the CPU Monitor was active.

While the response time and emulation phases are independent, the reports produced are complimentary -- there are several auxiliary reports generated by the response time phase useful for determining the distribution of CPU resources across the entire system and also within TSS.

An optional third pass is an investigative and debug tool for the more technical user. It provides a formatted dump for examining TSS transaction content, sequencing and anomaly.

TEARS produces tabular reports, histograms, and logarithmic bar charts as tools for analysis. Of the three types, the logarithmic bar chart is new to GMF reduction programs - it is a variation of the standard plot seen in other sections of this manual. Most of the bar charts described in following subsections represent, in two ways, the duration of some frequently occurring process. The chart is ordered by clock time and is arranged vertically on the page. The time interval between print lines is under user control with a granularity of one second. For each print line, the left half of the chart is a histogram (or interval-o-gram) showing how many such processes were observed in each of several duration lengths (buckets) during the observation interval. If the process is attributable to a TSS user, the user with the longest duration is identified under the column heading "LINE". The right half of the chart prints the minimum,

* Independently, the CAM Monitor and associated data reduction package can identify periods of poor response.

the average, and the maximum duration as horizontal bars on a logarithmic scale. The usual way of printing a multivariable bar chart is to lay the bars side by side:

	minimum	mmm
time	average	aaaaaa
	maximum	MMMMMMMM

To make the charts compact, we have overlaid the average bar with the minimum bar and then that composite bar overlays the maximum bar:

time	mmmaaaMMM
------	-----------

If 2 or more bars are of equal length, a digit "2" or "3" is printed in lieu of the characteristic bar symbol. (A "2" is used if 2 bars are of equal length; "3", if all 3 bars are of equal length.) If one or more bars fails to cross the left margin (axis), the axis symbol, "I", is replaced with an asterisk. If one or more bars would break the right margin, an asterisk is printed in lieu of the bar symbol.

For each observation interval (print interval), a process is counted if (1) it terminates during the observation interval, or (2) it is ongoing at the end of the interval. In either case, the duration is measured from the beginning of the process, not from the beginning of the observation interval. Therefore, a process which spans two or more observation intervals will be counted once in each observation interval, and its reported duration may exceed the observation period.

These bar charts are intended as visual displays. It is patterns and trends and order-of-magnitude changes that are important, not the ability to convert a bar length to a numeric with accuracy (the line printing mechanism forbids accuracy anyway). The default method of displaying the y-axis shows numeric values at each order-of-magnitude only:

1.OEO1	1.OEO2
Y-----	Y-----

By user option (option (logscale)), intermediate values can be displayed:

1.OEO1	1.OEO2
Y-----2-----3-----4-----5-----6--7-8-9-Y	

15.4 Data Reduction Methodology

TEARS is configured to correlate up to 50 concurrent TSS users (real and deferred) with user-generated traces. If the GMC data tape shows a TSS user load greater than 50, TEARS will terminate the reduction pass and print all accumulated reports to that point. To increase (or decrease) the 50-user limit, edit, then recompile the source file B29IDPX0/SOURCE/TEARS as follows:

CASE

B RVS:/USTSUP = 50/*:/USTSUP = NN/

where NN is the new concurrency limit. It is fruitless to set NN beyond 100, because the TSS Monitor will generate no traces with a TSS user load that large.

TEARS, like most other GMF packages, uses the system controller clock to measure the passage of time. Precision of the clock is 1 microsecond, and its capacity is 35 bits. Unlike the other GMF packages, no provision exists in TEARS to account for clock overflow. TEARS is thereby limited to reducing GMC sessions shorter than approximately 9.5 hours. On reaching the time limit, TEARS will terminate the reduction pass gracefully.

15.4.1 Response Pass Reduction Methodology. This pass performs an overview of TSS response during the data collection period. The reports generated give an indication of the periods of time when response, as observed by the user, was poor.

To accomplish this, the response pass keeps track of the following information:

- Start and stop of all terminal I/Os
- Start and stop of all user subdispatches
- Start and stop of all disk I/Os
- Requests for core by subsystems
- Changes of each user's program stack
- Special situations (JDAC, CONN, etc.).

To produce the response time logplots, the program measures the amount of time from the completion of a terminal I/O until the initiation of a subsequent output transmission. This value, the line-idle period, is used as a datum for the "Response Time For All Users" logplot, provided none of the following situations have occurred:

If the user's subsystem executes a DRL WAKE, it is assumed that the user knows that the line idle is being extended (e.g. DJST). If the user performed a line-switch to another program, the line will be idle as far as TSS is concerned, but not used in the logplots. If the user performed a non-graceful logoff (e.g. \$*\$D, or CTRL-C), the UST which was set for reconnect will not contribute to the logplot.

Additionally, each datum used in the All-User logplot will be used in one of the other two response time logplots. The distribution of data between the two logplots is based on the presence of a request-for-core trace (TSS subtypes 77-78) during the line-idle period. This produces a segregation of responses for users in build mode or with program-in-core versus users requiring a program call or normal swap in. Comparison of these reports will indicate when poor response may be due to core contention (provided sufficient users are logged on to produce enough plotting data). This

problem would be indicated by a degradation on the "Response Time For Users With Core Request During Line Idle" logplot without a corresponding one on the "Response Time For Users Not Requesting More Core" logplot.

These logplots can be distorted by a user who performs an extensive amount of disk I/O or CPU work without doing any terminal I/O. The effect would show up as a long period of line idle because the reduction program could not invalidate that time datum. As a run-time option, the amount of time spent doing disk I/O and the amount of processor time can be subtracted from each datum, thus reducing, but not eliminating, the erroneous effect.

To produce the subdispatch logplots, the program measures the amount of time from the point where a user is inserted into the subdispatch queue, to the point where the user is processed as a fault entry after receiving a subdispatch. The logplot "Total Time In Subdispatch Queue" reflects this time value. The logplot "Time in Subdispatch Queue Waiting Service" is the total time value minus the actual processor time used by the subsystem. The latter value has been separated out to produce the logplot "Processor Time in Subdispatch."

15.4.2 Emulation Reduction Methodology. This, the second phase of TEARS, simulates the TSS process in a limited fashion. Each user is placed in one of several defined states (table 15-1) corresponding to the manner in which TSS processes a user. For each user, TEARS maintains a state-stack; pushing, popping, or clearing of the state-stack is driven by a complex algorithm. The nature of the current trace for the user, the user's Time Type (a TSS classification), and his current and previous states interact in deciding which state-stack operation will be performed. Several states are considered vulnerable for a user -- if he remains in one of these states too long, a service for him is being delayed or denied. In a similar fashion, the TSS Executive may be denied CPU access thus delaying all users. These separate events, user delays in vulnerable states, and TSS Executive delays, are examined and reported on together so that the effect of one on the other may be seen.

As each user transits from state to state, several statistics are updated. They indicate possible bottlenecks in CPU processing, I/O, or memory allocation from minute to minute. The statistics are taken across all users and provide a functional representation of TSS bottlenecks. (The CPU reports are created in the Response phase described in section 15.5.1.4, but find application in the Emulation phase).

Concurrently with the simulation, each trace is examined as an indication of TSS trouble. Just by their occurrence, several traces indicate TSS or GCOS trouble. Over 40 conditions, many concerned with derail errors, are taken as exceptions and are items for reporting.

Table 15-1. Emulation Phase User States

<u># Mnemonic</u>	<u>Description</u>	<u>Vulnerable?</u>
1 GWAKE	In Gewake	
2 LIDLE	Line is idle	
3 WTSUB	Eligible for subdispatch	
4 RECON	In reconnect mode	
5 SWPOU	Swapping out	Yes
6 SWPIN	Swapping in	Yes
7 DSKIO	Doing disk I/O	Yes
8 SUBDS	In subdispatch queue	Yes
9 WTMEM	Waiting for memory	Yes
10 DRLSR	Waiting for derail service	Yes
11 NOTSS	Non-TSS process	
12 FMSE	Waiting FMS service	Yes
13 WTSME	Waiting SMC service	Yes
14 LBUSY	Line is busy	
15 IDLME	Idle in memory	
16 UNKWN	Unknown (initial state)	
17 LOGOF	Logged off (final state)	

15.5 Output

Each of the three reduction phases (passes) provides independent but complementary reports. Each report begins with a banner line identifying the TEARS program, the version, and the version date. The banner is usually followed by a line identifying the data source -- this includes system identification, time of day, day of week, date (year-month-day) and reel number of the lead-off GMC data tape. The day of the week and date are continually updated and reflect the correct GMC start time or, when user selected, the time datum of the first collectable trace within a timeframe. Following the data source line is the report title and report number.

The reports generated by the TEARS reduction passes fall into three categories: summary, interval driven, and event driven. Summary reports are generated at the end of a timeframe or at the end of the pass (e.g., Histograms). Interval driven reports are generated at the end of a short time period (defaulting to 60 or 20 seconds) repeatedly during a timeframe (e.g., Logplots). Event driven reports contain data generated when certain traces or conditions are encountered during the processing of the collection tape. The information is organized as a chronological listing.

One report, the Execution Summary (figure 15-1), is common to all three reduction phases; it shows what TEARS found on the data files. Included in this information is the following:

- o A list of the monitors in execution during the GMC data collection.
- o The time, date, tape number and system controller clock at the beginning of collection.
- o An echo of the user input options, and the program's interpretation of them.
- o If the timeframe option is used, a report of when the various timeframes were reached.
- o If the collection period exceeded 9+ hours, an indication that a 35-bit internal clock overflowed (TEARS will then terminate the pass).
- o The time, date, tape number and system controller clock at the end of data collection, but only if a termination trace is processed (not currently implemented).
- o When the NEW option is user selected, the items above are repeated on a new page.

All remaining reports are generated by the individual reduction passes.

The reports that the Response pass produces are:

Trace driven reports:

- o TSS Reduction Event Log (figure 15-2) (no ID #).

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

ACTIVE MONITORS: **** CPU ***** TSS ****

DATA COLLECTION BEGAN WITH TAPE 1469 AT 0818:33.201 FRI 82-05-21,
 RSCRA/Q WERE 00000000031 460323216133

CARD INPUT IS INTERPRETED AS FOLLOWS:

1:OPTION(-SUBDISPATCH) OPTION(-DISKIO)
 WILL DO -SUBDISPATCH
 WILL DO -DISKIO

2:CPRI(20)
 PERIODIC REPORT REPETITION PERIOD IS NOMINALLY 20 SECONDS.

3:NEW(1469)
 TAPE 1469 WILL LEAD OFF A FOLLOWING REDUCTION.
 *--> BREAKPOINT AT 0818:54.01

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

ACTIVE MONITORS: **** CPU ***** TSS ****

DATA COLLECTION BEGAN WITH TAPE 1469 AT 0818:33.201 FRI 82-05-21,
 RSCRA/Q WERE 00000000031 460323216133

CARD INPUT IS INTERPRETED AS FOLLOWS:

1:ON(FORMATTING)
 FORMATTED DUMP TURNED ON .

Figure 15-1. Execution Summary Report

TIME SHARING REDUCTION EVENT LOG

```

***** MAIN INITIALIZED ***** 20.00 SET. PLOT INTERVAL *****
0 *** TIMEFRAME ***** BEIN *****
1- 90- ***** TRACE ON *****
5- 90- ***** 4 USERS *****
0818:33.20 7140- 71- ** UST AREA INCREASE. 5 USERS.
0818:43.51 7184- 12- (CT)-NEM LOGON HIS5010
0818:54.01 8785- 12- (CT)-NEM LOGON C433
0823:21.44 16814- 12- (CT)-NEM LOGON C722
0824:06.74 17155- 71- ** UST AREA INCREASE. 7 USERS.
0826:45.04 21220- 12- (CT)-NEM LOGON NSSIB
0828:44.92 38142- 12- (CT)-NEM LOGON PRC118
0833:41.42 41818- 44- (CT)-JDAC- LINE SWITCH TO INQUIRY NAME VIDEO
0835:48.38 66305- 12- (CT)-TERM LOGOFF PRC118 ON AT 834.41
0842:34.15 67126- 12- (CT)-NEM LOGON SSLSM
0843:09.43 68551- 12- (CT)-NEM LOGON PRC101
0844:47.87 71175- 12- (CT)-TERM LOGOFF PRC101
0844:18.47 72586- 98- ** UST AREA INCREASE. 9 USERS.
0844:45.77 72663- 12- (CT)-NEM LOGON TRCN11
0844:47.35 74961- 12- (CT)-NEM LOGON HIS
0845:15.51 75536- 12- (CT)-NEM LOGON CSC
0845:20.80 75739- 71- ** UST AREA INCREASE. 13 USERS.
0845:21.20 75789- 12- (CT)-NEM LOGON PRC101
0845:21.31 75789- 12- (CT)-NEM LOGON PRC101
0847:19.78 96856- 22- (CT)-EDTA- LINE SET FOR RECONNECT
0847:19.78 96856- 12- (CT)-TERM LOGOFF HIS ON AT 845.15
0847:22.77 96880- 22- (CT)-CML- LINE SET FOR RECONNECT
0847:22.78 96880- 22- (CT)-DUST- LINE SET FOR RECONNECT
0848:36.66 100204- 12- (CT)-NEM LOGON CSC ON AT 848.35
0848:59.64 104802- 12- (CT)-TERM LOGOFF PRC101
0849:00.78 105160- 12- (CT)-NEM LOGON PRC101
0849:15.98 105889- 58- (CT)-CORN- CORN COMMAND TO -72-
0849:16.21 105968- 12- (CT)-TERM LOGOFF PRC101 ON AT 849.01
0849:54.15 107675- 12- (CT)-NEM LOGON ACTMIS ON AT 816.23
0850:19.96 108868- 12- (CT)-TERM LOGOFF PRC104 ON AT 824.07
0850:20.14 108893- 12- (CT)-TERM LOGOFF C433
0850:30.26 109760- 12- (CT)-NEM LOGON CSC7592 ON AT 828.45
0851:44.14 113669- 12- (CT)-TERM LOGOFF NSSIB ON AT 823.21
0852:21.28 114942- 12- (CT)-TERM LOGOFF HIS5010
0852:53.32 116852- 12- (CT)-NEM LOGON C432
0854:03.46 121808- 12- (CT)-NEM LOGON PRC118
0854:05.48 122591- 12- (CT)-NEM LOGON PRC105
0855:19.36 130701- 12- (CT)-NEM LOGON CSC8121
0856:50.36 143098- 22- (CT)-NEM- LINE SET FOR RECONNECT ON AT 843.09
0859:48.10 179626- 12- (CT)-TERM LOGOFF SSLSM ON AT 859.98
0859:59.61 180159- 12- (CT)-NEM LOGON CSC ON AT 852.53
0900:14.27 184775- 12- (CT)-TERM LOGOFF C432
0900:53.82 188034- 12- (CT)-NEM ** UST AREA DECREASE. 12 USERS.
0900:53.82 188035- 99- ** UST AREA INCREASE. 13 USERS.
0903:06.37 222036- 22- (CT)-JDAC-
0903:06.72 222107- 71- ** UST AREA INCREASE. 13 USERS.
0903:10.12 222877- 12- (CT)-NEM LOGON HASB
0903:21.18 224095- 12- (CT)-CORN- CORN COMMAND TO -DI-
0903:21.43 224170- 12- (CT)-TERM LOGOFF HASB ON AT 906.06
0903:03.30 234008- 12- (CT)-TERM LOGOFF PRC118 ON AT 854.05
0912:45.40 272575- 12- (CT)-TERM LOGOFF HASB ON AT 814.41
0913:16.10 281724- 12- (CT)-NEM LOGON HIS
0913:16.10 281724- 12- (CT)-TERM LOGOFF CSC7592 ON AT 850.30
0921:31.57 319884- 12- (CT)-TERM LOGOFF PRC101 ON AT 845.21
0922:03.71 325271- 12- (CT)-NEM LOGON PRC101
0922:03.71 325271- 12- (CT)-TERM LOGOFF HIS
0923:30.73 400020- 12- (CT)-NEM LOGON TRCN11
0929:41.40 412835- 22- (CT)-EDTA- LINE SET FOR RECONNECT

```

Figure 15-2. Timesharing Reduction Event Log Report

Elapsed-time driven reports:

- o Response Time For All Users (figure 15-3) (ID #4)
- o CPU Times Allocated To TSI-EXEC and TSI-SUBDISPATCH (figure 15-4) (no ID #) (subordinate to the above report)
- o Response Time For Users Not Requesting More Core (figure 15-5) (ID #6)
- o Response Time For Users With Core Request During Line Idle (figure 15-6) (ID #5)
- o Total Time In Subdispatch Queue (figure 15-7) (ID #10)
- o Time In Subdispatch Queue Waiting Service (figure 15-8) (ID #11)
- o Processor Time In Subdispatch (figure 15-9) (ID #12).

Summary reports:

- o TSS Subtraces Encountered (figure 15-10 parts 1 and 2) (ID #1 & #2)
- o Derails Executed By Users (figure 15-11 parts 1 and 2) (ID #7 & #8).

The reports that the Emulation pass produces are:

Trace driven reports:

- o Exception Message Report (Tabular) (figure 15-12) (no ID #) and
- o Users Map (Tabular) (figure 15-13) (ID #3) (subordinate to above report)
- o TSS Delay & User Delay Report (Tabular) (figure 15-14) (no ID #) and
- o TSS Core Map (Tabular) (figure 15-15) (no ID #) (subordinate to above report)
- o Error Message Report (Tabular) (figure 15-16) (no ID #).

Elapsed-time driven reports:


- o Intertrace Gap - Duration (Tabular and Logplot) (figure 15-17) (ID #13)
- o Memory Activity (Tabular) (figure 15-18) (no ID #) (subordinate to above report)
- o Eligibles For Subdispatch - Duration (Tabular and Logplot) (figure 15-19) (ID #14)
- o In Subdispatch - Duration (Tabular and Logplot) (figure 15-20) (ID #15)
- o Eligible For & In Subdispatch - Duration (Tabular and Logplot) (figure 15-21) (ID #16)
- o User Swaps - Swap Rate (Logplot) (figure 15-22) (ID #17)
- o User Swaps - Swap Amount (Tabular and Logplot) (figure 15-23) (ID #18)
- o User Swaps - Duration (Tabular and Logplot) (figure 15-24) (ID #19)
- o User I/Os - Duration (Tabular and Logplot) (figure 15-25) (ID #20)
- o In System Master Catalog Wait - Duration (Tabular and Logplot) (figure 15-26) (ID #21).

Summary reports:

- o Intertrace Time (Adjusted for TSS GEWAKES) (Histogram) (figure 15-27) (ID #9).

3. REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82
 DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL TAP# 1464

RESPONSE TIME FOR ALL USERS
 -- MIN /- AVG -- MAX
 PLOT- 1



CURVE MINIMUM MAXIMUM
 MIN 2.64E-01 4.34E-01
 AVG 2.19E-02 1.21E-02
 MAX 1.03E-01 1.00E-03

Figure 15-3. Response Time For All Users Report

REPORT- 1

CPU TIMES ALLOCATED TO TSI-EXEC AND TSI-SUBDISPATCH


PLOT TIME REFERENCE	FIRST TFO TIME	LAST TFO TIME	NUMBER OF TFOs	CPU SECONDS ATTRIBUTABLE TO THE FOLLOWING				TOTAL
				TSI-EXEC	TSI-SUBSP	IDLE-TIME	OTHER-CPU	
0819:03.5	0818:43.5	0818:43.5	1	0.25	0.04	0.73	0.01	0.30
0819:23.5	0819:17.0	0819:17.0	2	0.73	0.04	65.73	0.70	67.21
0819:43.5	0819:23.7	0819:23.7	1	0.40	0.01	12.58	0.34	13.33
0820:03.5	0819:48.3	0819:48.3	1	0.35	0.02	48.37	0.39	49.12
0820:23.5	0819:48.3	0820:12.7	1	0.19	0.02	48.27	0.36	49.01
0820:43.5	0820:12.7	0820:30.7	1	0.21	0.01	35.33	0.47	35.99
0821:03.5	0820:30.7	0820:55.5	1	0.74	0.01	48.92	0.42	49.52
0821:23.5	0820:55.5	0821:19.2	2	0.65	0.24	45.72	0.66	47.35
0821:43.5	0821:19.2	0821:43.1	1	0.96	0.09	46.53	0.66	47.93
0822:03.5	0821:43.1	0822:03.2	11	0.46	0.07	32.74	6.52	40.33
0822:23.5	0822:03.2	0822:12.1	4	0.34	0.02	13.86	3.55	17.89
0822:43.5	0822:12.1	0822:38.1	1	0.73	0.02	50.69	0.38	51.43
0823:03.5	0822:38.1	0823:21.4	4	1.17	0.21	84.80	1.47	87.22
0823:23.5	0823:21.4	0823:42.1	3	0.38	0.24	39.09	1.07	41.57
0824:03.5	0823:42.1	0824:02.6	1	0.16	0.02	79.33	1.17	80.88
0824:23.5	0824:02.6	0824:07.0	1	0.71	0.00	8.28	0.36	8.81
0824:43.5	0824:07.0	0824:32.1	1	1.11	0.67	48.24	0.62	50.24
0825:03.5	0824:32.1	0825:02.9	3	0.55	1.05	58.51	0.96	61.63
0825:23.5	0825:02.9	0825:20.1	2	0.36	0.45	32.88	0.66	34.54
0825:43.5	0825:20.1	0825:39.4	2	1.74	0.04	36.88	1.42	38.70
0826:03.5	0825:39.4	0826:01.4	4	0.34	5.53	44.47	2.30	44.04
0826:23.5	0826:01.4	0826:20.8	4	1.01	0.04	36.87	1.41	38.83
0826:43.5	0826:20.8	0826:42.2	6	2.27	2.52	37.46	1.87	42.86
0827:03.5	0826:42.2	0827:00.4	2	0.66	3.32	27.20	1.69	34.48
0827:23.5	0827:00.4	0827:14.8	6	0.24	0.01	28.53	0.61	30.75
0827:43.5	0827:14.8	0827:40.4	1	0.25	0.01	50.59	0.41	51.26
0828:03.5	0827:40.4	0828:06.5	1	0.65	0.19	51.47	0.38	52.29
0828:23.5	0828:06.5	0828:39.1	4	0.83	0.27	62.86	1.47	65.23
0828:43.5	0828:39.1	0829:01.1	3	0.81	0.07	42.11	1.04	44.05
0829:03.5	0829:01.1	0829:23.1	2	1.19	0.05	42.59	0.65	44.08
0829:23.5	0829:23.1	0829:40.8	13	0.83	0.27	31.66	2.23	35.25
0830:03.5	0829:40.8	0830:02.4	3	0.76	0.03	32.85	0.81	34.52
0830:23.5	0830:02.4	0830:23.1	5	1.16	0.03	39.29	1.35	41.42
0830:43.5	0830:23.1	0830:41.2	10	0.73	0.96	32.05	2.43	36.17
0831:03.5	0830:41.2	0831:03.3	9	1.54	0.05	27.82	5.30	34.32
0831:23.5	0831:03.3	0831:21.0	3	0.56	0.03	27.22	6.55	35.61
0831:43.5	0831:21.0	0831:44.9	3	0.87	0.29	45.74	1.40	47.73
0832:03.5	0831:44.9	0832:01.5	3	0.87	0.29	31.10	1.11	33.37

Figure 15-4. CPR Times Allocated to Exec and Subdispatch Report

PLOT-2

RESPONSE TIME FOR USERS NOT REQUESTING MORE CORE

-- MIN /- AVG +- MAX



CURVE MINIMUM MAXIMUM
MIN -2.648E-01 9.5021E 00
AVG 5.8324E-03 2.9651E 02
MAX 1.5157E-02 1.0098E 03

Figure 15-5. Response Time For Users Not Requesting More Core Plot

PLOT- 3

RESPONSE TIME FOR USERS WITH CORE REQUEST DURING LINE IDLE

-- MIN /- AVG +- MAX




Figure 15-6. Response Time For Users With Core Requests During Line Idle Plot

PICT-4




Figure 15-7. Total Time In Subdispatch Queue Plot


... REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSC01 AT 0018:43.515 PM 82-05-21. INITIAL CMC TAPE #1469

PLOT- 5

TIME IN SUBDISPATCH QUEUE WAITING SERVICE

-- MIN -- AVG -- MAX



CURVE
MIN 2.540E-01
AVG 8.0279E 01
MAX 9.7754E 01


Figure 15-8. Total Time In Subdispatch Queue Waiting Service Plot

TSS REMISSION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82
 DISTRIBUTION COLLATED ON SYSTEM ISCT AT 0818:43.515 FRI 82-05-21. INITIAL CMC TAPE #1464

PLOT- 6

PROCESSOR TIME IN SUBDISPATCH

-- MIN /- AVG -- MAX



CURVE
 MIN 1.5600E-01
 AVG 2.2800E-01
 MAX 3.4300E-01

MAXIMUM
 1.8700E-01
 1.6400E-01
 4.0156E-01

Figure 15-9. Processor Time In Subdispatch Plot

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 PRI 82-05-21. INITIAL CMC TAPE #1469

TSS SUBTRACES ENCOUNTERED (LOW HALF)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS SUBTRC	PERCENT OF OCCURRENCE	HISTOGRAM
70	70	0.006	0.006	1-	0	LOG-ON COMPLETE
187	257	0.006	0.006	2-	1	SNAP SPACE DENIED
719	271	0.004	0.004	3-	1	PAT DENIAL
719	1696	0.153	0.066	4-	1	GEPSIE DEALLOCATE
1564	1702	0.153	0.001	5-	1	DEALOT COMPLETE
2207	3266	0.208	0.202	6-	1	ERROR MESSAGE
928	6401	0.595	0.085	7-	1	SYSTEM PRIMITIVE
1801	8202	0.750	0.165	8-	1	ENTER BUILD MODE
117	8319	0.761	0.011	9-	1	EXEC PRIMITIVE
53	8372	0.766	0.005	10-	1	LOG-OFF
1237	9609	0.879	0.113	11-	1	SYSTEM PRIMITIVE
9705	19314	1.766	0.887	12-	1	LOG-OFF
0	19314	1.766	0.000	13-	1	COMMAND RECEIVED
53	19367	1.771	0.005	14-	1	PERIODIC CHECK
5160	24527	2.245	0.472	15-	1	GEPSIE NO USERS
1560	18123	16.209	0.088	16-	1	BREAK/DISCONNECT
1081	18220	16.209	0.009	17-	1	GEPSIE S/D DONE
58	18266	16.703	0.005	18-	1	GEPSIE S/O BUSY
14	18267	16.704	0.001	19-	1	APR RTO I/C
95117	27772	25.401	8.698	20-	1	BUILD I/O COMPLETE
31902	30964	26.319	2.917	21-	1	RECONNECT MODE
29583	33927	31.024	2.705	22-	1	REQUEST FILE I/O
0	33927	31.024	0.000	23-	1	DISK I/O COMPLETE
1755	341032	31.184	0.160	24-	1	DRI DENIAL
187	341219	31.201	0.017	25-	1	ISSUE WME GEPSIE
1755	342974	31.362	0.160	26-	1	BAD FILACT PARAM
52	343026	31.367	0.005	27-	1	GEPSIE/27 COMPLETE
36	343062	31.370	0.003	28-	1	DELAY 200 MS
70	343132	31.376	0.006	29-	1	DELAY 2 SEC
20	343152	31.378	0.002	30-	1	STORE USERID
0	343152	31.378	0.000	31-	1	PASIST -1 COMP
91	343242	31.386	0.008	32-	1	DRI MORLINK ERROR
91	343274	31.392	0.008	33-	1	SPAWN I/O COMPLETE
7	343314	31.399	0.001	34-	1	BATCH JOB
91	343345	31.404	0.001	35-	1	DRI SPAN ERROR
376	343527	31.412	0.008	36-	1	WRITE ABST
928	343531	31.447	0.034	37-	1	REPT. I/O COMPLETE
4	344511	31.532	0.085	38-	1	REPT. I/O DENIAL
15	344850	31.533	0.001	39-	1	LOAD DRI RESTOR
15	344865	31.535	0.001	40-	1	RESTOR I/O RESTOR
54	344866	31.538	0.003	41-	1	RESTOR I/O RESTOR
2728	344899	31.538	0.003	42-	1	START LINE SWITCH
0	344899	31.538	0.000	43-	1	FILE GROW .RES19
0	34627	31.787	0.249	44-	1	GROW COMPLETE
0	34627	31.787	0.000	45-	1	CONSOLE I/O
0	34627	31.787	0.000	46-	1	DRI JOUT RETURN
0	34627	31.787	0.000	47-	1	DRI JOUT BUSY
0	34627	31.787	0.000	48-	1	START GMAKE
0	34627	31.787	0.000	49-	1	*J WRITE - TASK
0	34627	31.787	0.000	50-	1	START DRI TASK
0	34627	31.787	0.000	51-	1	*J READ - TASK
0	34627	31.787	0.000	52-	1	DRI SAVE COMPLETE
0	34627	31.787	0.000	53-	1	
0	34627	31.787	0.000	54-	1	

1094606 ENTRIES TOTAL 745976 OUT OF RANGE

Figure 15-10. TSS Subtraces Encountered Histogram (Part 1 of 2)


INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS SUBTRC	PERCENT OF OCCURRENCE	HISTOGRAM
					00 05 10 15 20 25 30 35 40 45 50	
347670	347630	31.787	31.787	1-	XXXXXXXXXXXXXXXXXXXX	IDL SAVE COMPLETE
15	347645	0.001	31.789	54	XXXXXXXXXXXXXXXXXXXX	IDS ATTRI .MFS3
15	347660	0.001	31.790	55	XXXXXXXXXXXXXXXXXXXX	IDS ATTRI STATUS
0	347660	0.	31.790	56	XXXXXXXXXXXXXXXXXXXX	DRL T.SYST DENIAL
3	347663	0.	31.791	57	XXXXXXXXXXXXXXXXXXXX	LIST SWITCH T.CONN
0	347663	0.	31.791	58	XXXXXXXXXXXXXXXXXXXX	TAP= I/O COMPLETE
68443	416106	38.049	6.258	59	XXXXXXXXXXXXXXXXXXXX	ALLOCATOR SERVICES
29546	675652	61.782	23.733	60	XXXXXXXXXXXXXXXXXXXX	ENTER PROC ALLOC
34996	710648	64.982	3.200	61	XXXXXXXXXXXXXXXXXXXX	ENTER MEMORY ALLOC
1206	721054	65.934	0.841	62	XXXXXXXXXXXXXXXXXXXX	ENTER SNAP DECIS
9200	721059	65.934	0.000	63	XXXXXXXXXXXXXXXXXXXX	NEXT PRIMITIVE IS
5	721062	65.934	0.000	64	XXXXXXXXXXXXXXXXXXXX	CONSIDER TSS INCR
3	721071	65.935	0.001	65	XXXXXXXXXXXXXXXXXXXX	INITIATE SIZE INCR
9	721072	65.935	0.000	66	XXXXXXXXXXXXXXXXXXXX	SET UP FENCE
1	721084	65.936	0.001	67	XXXXXXXXXXXXXXXXXXXX	MEMORY RESERVE
12	721084	65.936	0.	68	XXXXXXXXXXXXXXXXXXXX	FORCE SNAP
0	721084	65.937	0.001	69	XXXXXXXXXXXXXXXXXXXX	TERMINATE SNAP
10	721094	65.937	0.001	70	XXXXXXXXXXXXXXXXXXXX	TSS1 LIST AREA +1K
7	721101	65.938	0.000	71	XXXXXXXXXXXXXXXXXXXX	REMOVE MEMORY
5	721106	65.938	0.000	72	XXXXXXXXXXXXXXXXXXXX	REMOVE SUCCESSFUL
5	721109	65.939	0.000	73	XXXXXXXXXXXXXXXXXXXX	MEMORY RELEASE
137	721246	65.951	0.013	74	XXXXXXXXXXXXXXXXXXXX	REMOVE REQUESTED
12	721258	65.952	0.001	75	XXXXXXXXXXXXXXXXXXXX	CANCEL CRUI
9709	730667	66.840	0.888	76	XXXXXXXXXXXXXXXXXXXX	ATTEMPT NEW ALLOC
23932	762819	69.028	2.188	77	XXXXXXXXXXXXXXXXXXXX	NEW MAP CHANGE
48418	810635	74.116	4.427	78	XXXXXXXXXXXXXXXXXXXX	MAP PROGRAM
2646	813181	74.358	0.242	79	XXXXXXXXXXXXXXXXXXXX	TIME TYPES
896	814077	74.440	0.082	80	XXXXXXXXXXXXXXXXXXXX	REQUEST REM
62	814139	74.445	0.002	81	XXXXXXXXXXXXXXXXXXXX	SCAN COMMAND LIST
3124	817263	74.731	0.286	82	XXXXXXXXXXXXXXXXXXXX	SCAN REQUESTED
632	817895	74.789	0.058	83	XXXXXXXXXXXXXXXXXXXX	THRU COMPLETE
17	817912	74.790	0.002	84	XXXXXXXXXXXXXXXXXXXX	POPUP TO CALLS
11195	829107	75.814	1.024	85	XXXXXXXXXXXXXXXXXXXX	TPB COMPLETE
97904	927011	81.763	8.953	86	XXXXXXXXXXXXXXXXXXXX	TERM I/O COMPLETE
11909	938920	85.855	1.086	87	XXXXXXXXXXXXXXXXXXXX	S/D FAULT QUEUE
5	938925	85.856	0.002	88	XXXXXXXXXXXXXXXXXXXX	PROCESS .CRPSQ
97905	1065370	84.815	0.006	89	XXXXXXXXXXXXXXXXXXXX	TRACE ON/OFF
10	1036000	84.815	0.002	90	XXXXXXXXXXXXXXXXXXXX	MAKE S/D ENTRY
0	1036000	84.815	0.	91	XXXXXXXXXXXXXXXXXXXX	LOG-ON REQUEST
67	1036357	84.821	0.006	92	XXXXXXXXXXXXXXXXXXXX	BAD LINE STATUS
0	1036370	84.821	0.000	93	XXXXXXXXXXXXXXXXXXXX	CHECK VIP TERM TYP
73	1037043	84.828	0.007	94	XXXXXXXXXXXXXXXXXXXX	REJECT USER
4	1037047	84.828	0.000	95	XXXXXXXXXXXXXXXXXXXX	LIST COMPRESSION
5	1037052	84.829	0.000	96	XXXXXXXXXXXXXXXXXXXX	TSSN LIST AREA +1K
18549	1055601	96.525	1.696	97	XXXXXXXXXXXXXXXXXXXX	LIST AREA -1K
93	1055694	96.533	0.009	98	XXXXXXXXXXXXXXXXXXXX	TERMINAL I/O
67	1055761	96.539	0.006	99	XXXXXXXXXXXXXXXXXXXX	CRUN \$\$\$ FUNC
268	1056029	96.564	0.025	100	XXXXXXXXXXXXXXXXXXXX	REMOTE INQUIRY
37571	1093606	100.000	3.436	101	XXXXXXXXXXXXXXXXXXXX	VIP INPUT COMPLETE
				102	XXXXXXXXXXXXXXXXXXXX	DRL COMPLETE
				103	XXXXXXXXXXXXXXXXXXXX	
				104	XXXXXXXXXXXXXXXXXXXX	
				105	XXXXXXXXXXXXXXXXXXXX	
				106	XXXXXXXXXXXXXXXXXXXX	
				107	XXXXXXXXXXXXXXXXXXXX	
				108	XXXXXXXXXXXXXXXXXXXX	
				109	XXXXXXXXXXXXXXXXXXXX	
				110	XXXXXXXXXXXXXXXXXXXX	
				111	XXXXXXXXXXXXXXXXXXXX	
				112	XXXXXXXXXXXXXXXXXXXX	
				113	XXXXXXXXXXXXXXXXXXXX	
				114	XXXXXXXXXXXXXXXXXXXX	
				115	XXXXXXXXXXXXXXXXXXXX	
				116	XXXXXXXXXXXXXXXXXXXX	
				117	XXXXXXXXXXXXXXXXXXXX	
				118	XXXXXXXXXXXXXXXXXXXX	
				119	XXXXXXXXXXXXXXXXXXXX	
				120	XXXXXXXXXXXXXXXXXXXX	

1093606 ENTRIES TOTAL

0 OUT OF RANGE

Figure 15-10. (Part 2 of 2)

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL CMC TAPE #1469

Figure 15-11. Derails Executed By User Histogram
(Part 1 of 2)

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82
 DISTRIBUTION COLLECTED ON SYSTEM BSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469




Figure 15-11. (Part 2 of 2)

EXCEPTION MESSAGE

[illegible]

Figure 15-12. Exception Message Report

USERS MAP

*****REASON FOR FOLLOWING MAP*****
 0818:56.604* 43 LINE <TN>, USERID PR0104 :: SWAP SPACE DENIED.

USR	USDT	LT	USER-ID	TIME-IN-STATE	TMTO	LAST-TRACE	ADERRAIL	SIZE	PROGRAM-STACK
1	065060	<TN>	PR0104	0	WTSUB	88 S/D FAULT QUEUE	0 ZERO	2KT	NONE CMDI
2	065530	<DI>	MASB	2591	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
3	065230	<CV>	PR0105	2591	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
4	062730	<TN>	0721	2591	UNKN	90 TRACE ON/OFF	0 ZERO	1K	CARD

*****REASON FOR FOLLOWING MAP*****
 0818:56.610* 49 LINE <TN>, USERID PR0104 :: CANNOT SWAP.

USR	USDT	LT	USER-ID	TIME-IN-STATE	TMTO	LAST-TRACE	ADERRAIL	SIZE	PROGRAM-STACK
1	065060	<TN>	PR0104	0	WTSUB	88 S/D FAULT QUEUE	27 PASIST	2KT	NONE CMDI
2	065530	<DI>	MASB	2597	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
3	065230	<CV>	PR0105	2597	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
4	062730	<TN>	0721	2597	UNKN	90 TRACE ON/OFF	0 ZERO	1K	CARD

*****REASON FOR FOLLOWING MAP*****
 0818:56.695* 55 LINE <TN>, USERID PR0104 :: I/O ERROR ON DEVICE.

USR	USDT	LT	USER-ID	TIME-IN-STATE	TMTO	LAST-TRACE	ADERRAIL	SIZE	PROGRAM-STACK
1	065060	<TN>	PR0104	0	WTSUB	29 CERSVE/21 COMPLETE	30 ELIAPT	2KT	NONE CMDI
2	065530	<DI>	MASB	2680	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
3	065230	<CV>	PR0105	2680	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
4	062730	<TN>	0721	2680	UNKN	90 TRACE ON/OFF	0 ZERO	1K	CARD

*****REASON FOR FOLLOWING MAP*****
 0818:56.933* 117 LINE <TN>, USERID PR0104 :: UNACCOUNTABLE ERROR 1.

USR	USDT	LT	USER-ID	TIME-IN-STATE	TMTO	LAST-TRACE	ADERRAIL	SIZE	PROGRAM-STACK
1	065060	<TN>	PR0104	0	WTSUB	88 S/D FAULT QUEUE	12 RETPIL	2KT	NONE CMDI
2	065530	<DI>	MASB	2920	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
3	065230	<CV>	PR0105	2920	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
4	062730	<TN>	0721	2920	UNKN	90 TRACE ON/OFF	0 ZERO	1K	CARD

*****REASON FOR FOLLOWING MAP*****
 0818:56.934* 121 LINE <TN>, USERID PR0104 :: CHECKSUM ERROR.

USR	USDT	LT	USER-ID	TIME-IN-STATE	TMTO	LAST-TRACE	ADERRAIL	SIZE	PROGRAM-STACK
1	065060	<TN>	PR0104	0	WTSUB	104 DRI COMPLETE	13 BEINIM	2KT	NONE CMDI
2	065530	<DI>	MASB	2921	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
3	065230	<CV>	PR0105	2921	UNKN	90 TRACE ON/OFF	0 ZERO	1K	NONE
4	062730	<TN>	0721	2921	UNKN	90 TRACE ON/OFF	0 ZERO	1K	CARD

Figure 15-13. Users Map

TSS DELAY REPORT --- USER DELAY REPORT									
TSS	RWT	DAY	PRV	CUR	ABT	DRI	--TSS DELAYED--		OR USERS SNAPSHOT (>)--
							PRV	DEL	
--- USER DELAYED ---									
--- DAY STATE T PRV DEFLT SIZE PROGRAM STACK ---									
--- DAY STATE T PRV DEFLT SIZE PROGRAM STACK ---									
0823:05.297	6371		61				1	DRISR	60 JOIT HK NONE JOIT
0823:05.299	6375		91				1	DRISR	60 JOIT HK NONE JOIT
0831:32.618	30344		61				3	DRISR	60 JOIT HK NONE JOIT
0831:32.620	30348		91				3	DRISR	60 JOIT HK NONE JOIT
0845:03.832	73475	1	18, 61	18			1	IDLNE	60 T.GOVY 3K TSBP
							1	LBUSY	78 KOUTN 34K RUNY
							41	LBUSY	100 RETURN 1K NONE
							33	LBUSY	78 KOUTN 26K NONE CHDI
							1	WTMDH	100 GWAKE 1K NONE CHDI
							12	LBUSY	100 KOUTN 5K NONE CHDI
							2	LBUSY	78 KOUTN 8K NONE JOIT
							546	NOTES	100 CGROUT 12K NONE JDAC
							5	LBUSY	100 KOUTN 11K CARD EUTX
							43	LBUSY	100 KOUTN 11K NONE EUTX
0846:03.890	84574		29*				1	PMSER	60 FILACT 5K NONE CHDI
0847:07.588	95759		61				2	DRISR	60 JOIT HK NONE GRUN JOIT
0847:07.567	95749		91				2	DRISR	100 JOIT HK NONE GRUN JOIT
0848:44.024*	100696		77				4	WTMDH	78 KOUTN 3K TSBP
0848:44.024*	100698		80				4	WTMDH	78 KOUTN 3K TSBP
0852:37.491	115212	6	23, 104				6	DRISR	87 DEFTL 34K MDQ CHDI
							63	GWAKE	100 GWAKE 1K EUTX FHN RUNY
							12	LBUSY	78 KOUTN 3K NONE CHDI
							0	WTMDH	104 DEFTL 34K MDQ CHDI
							343	LBUSY	100 RETURN 2K NONE
							2	WTMDH	78 GWAKE 1K NONE EUTX JRN
							95	LBUSY	78 KOUTN 26K NONE CHDI
							7	LBUSY	100 RETURN 1K NONE LIST
							1049	NOTES	100 CGROUT 12K NONE JDAC
							2	WTMDH	80 KOUTN 11K CARD EUTX
							28	LBUSY	100 KOUTN 11K NONE EUTX
0854:03.503	121745		29*				1	PMSER	60 FILACT 34K MDQ CHDI
0910:12.780	258364		100				2	DRISR	60 JOIT HK NONE CHDI JOIT
0910:12.781	258366		91				2	DRISR	100 JOIT HK NONE CHDI JOIT
0918:20.932	301132	2	23, 104				2	DRISR	87 DEFTL 34K MDQ CHDI
							0	WTMDH	104 DEFTL 34K MDQ CHDI
							139	LBUSY	100 KOUTN 3K RUNY

Figure 15-14. TSS Delay and Users Delay Report

TSS CORE MAP

```

*****REASON FOR FOLLOWING MAP*****
0848:44.024* 100696          71          : 11 9595 CS*          4 WTMDN  78 KOUTN  % TS8*
-----
USR BASE SIZE FREE  LTD USER-ID----- TSTAT  TPORE  TPCPU  DIOS  LSTDL  PROGRAM-STACK-----
H01E 31K  5K
-----
*****REASON FOR FOLLOWING MAP*****
0848:44.024* 100698          80          : 11 9595 CS*          4 WTMDN+  78 KOUTN  % TS8*
-----
USR BASE SIZE FREE  LTD USER-ID----- TSTAT  TPORE  TPCPU  DIOS  LSTDL  PROGRAM-STACK-----
11 31K  3K  55K  9595 CS*          0  4392  0  0 KOUTN  TS8*          071610 SMPIN  F  TIME TYPES
-----
*****REASON FOR FOLLOWING MAP*****
1031:42.147*104758          71          : 5 <TN> PR104  12 WTMDN  100 GWAKE  2K CARD EDTX DEXT
-----
USR BASE SIZE FREE  LTD USER-ID----- TSTAT  TPORE  TPCPU  DIOS  LSTDL  PROGRAM-STACK-----
H01E 36K  5K
25 41K  4K  45K <DI> MASA          0  16759  28  0 FILACT NONE MASE          064030 SUBRS  F  MAKE S/D ENTRY
-----
*****REASON FOR FOLLOWING MAP*****
1031:42.148*104756          80          : 5 <TN> PR104  12 WTMDN+  78 GWAKE  2K CARD EDTX DEXT
-----
USR BASE SIZE FREE  LTD USER-ID----- TSTAT  TPORE  TPCPU  DIOS  LSTDL  PROGRAM-STACK-----
5 36K  2K  3K <TN> PR104          0  17219  0  0 GWAKE  CARD EDTX DEXT          070010 SMPIN  F  TIME TYPES
25 41K  4K  45K <DI> MASA          2  16760  28  0 FILACT NONE MASE          064030 SUBRS  F  MAKE S/D ENTRY
-----
*****REASON FOR FOLLOWING MAP*****
1031:42.152*104756          80          : 12 <TN> WSSIB          10 WTMDN+  78 GWAKE  2K DEXT
-----
USR BASE SIZE FREE  LTD USER-ID----- TSTAT  TPORE  TPCPU  DIOS  LSTDL  PROGRAM-STACK-----
5 36K  2K  3K <TN> PR104          3  17222  0  0 GWAKE  CARD EDTX DEXT          070010 SMPIN  F  TIME TYPES
12 36K  2K  1K <TN> WSSIB          0  14541  12  20 GWAKE  DEXT          071570 SMPIN  F  TIME TYPES
25 41K  4K  45K <DI> MASA          5  16763  28  0 FILACT NONE MASE          064030 SUBRS  F  MAKE S/D ENTRY
-----
*****REASON FOR FOLLOWING MAP*****
1031:46.303*1047827          71          : 1 <TS> PR105          4 WTMDN  71 KOUTN  11K CARD EDTX
-----
USR BASE SIZE FREE  LTD USER-ID----- TSTAT  TPORE  TPCPU  DIOS  LSTDL  PROGRAM-STACK-----
H01E 36K  2K
12 36K  2K  3K <TN> WSSIB          1  928  1  2 DIO  DEXT          071570 SUBRS  F  MAKE S/D ENTRY
18 40K  2K  3K <TN> USERID UNKNW          2  646  0  0 CROUT NEW          073440 SUBRS  F  MAKE S/D ENTRY
25 42K  2K  4K <DI> MASA          0  144  0  0 RETURN NONE CMDI          064030 SUBRS  F  MAKE S/D ENTRY

```


Figure 15-15. TSS Core Map

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82
 DISTRIBUTION COLLECTED ON SYSTEM DISC1 AT 0818:33.201 FRI 82-05-21. INITIAL CMC TAPE #1469

TOD & RETNO	ERROR	MESSAGE	ERROR	MESSAGE
0850:19.630	108727	LINE CTK2, USERID C433	::	LAST KNOWN - WAIT NORM.
0920:36.929	3148308	LINE 7412, USERID DJE8C344AH	::	LAST KNOWN - WAIT FORCD.
				JUST REPORTED - OUT OF MEM.
				JUST REPORTED - OUT OF MEM.

PLOT-7

INTERTRACE GAP (DURATION) - ADJUSTED FOR TSS REMAKES
 -- MIN -- AVG -- MAX



CURVE
 MIN 0.
 AVG 4.3666E-01
 MAX 1.5940E 00

MINIMUM
 0.
 4.3666E-01
 1.5940E 00

MAXIMUM
 0.
 1.0942E 01
 1.2214E 04

Figure 15-17. Intertrace Gap Duration Plot

REPORT- 1

MEMORY ACTIVITY (TRACES 63, 65-75)


TOD	TRACE 63	TRACE 65	TRACE 66	TRACE 67	TRACE 68	TRACE 69	TRACE 70	TRACE 71	TRACE 72	TRACE 73	TRACE 74	TRACE 75
0818:54.012												
0819:14.012												
0819:34.012												
0819:54.012												
0820:14.012												
0820:34.012												
0820:54.012												
0821:14.012												
0821:34.012												
0821:54.012												
0822:14.012												
0822:34.012												
0822:54.012												
0823:14.012												
0823:34.012												
0823:54.012												
0824:14.012												
0824:34.012												
0824:54.012												
0825:14.012												
0825:34.012												
0825:54.012												
0826:14.012												
0826:34.012												
0826:54.012												
0827:14.012												
0827:34.012												
0827:54.012												
0828:14.012												
0828:34.012												
0828:54.012												
0829:14.012												
0829:34.012												
0829:54.012												
0830:14.012												
0830:34.012												
0830:54.012												
0831:14.012												
0831:34.012												
0831:54.012												
0832:14.012												
0832:34.012												
0832:54.012												

63-ENTER SWAP DEVIS
65-CONSIDER TSS INCR
66-INITIATE SIZE INCR
67-SET UP FENCE
68-MEMORY RESERVE
69-FORCE SWAP
70-TERMINATE SWAP
71-TSS1 LST AREA +1K
72-GENORE MEMORY
73-GENORE SUCCESSFUL
74-MEMORY RELEASE
75-GENORE REFUSED

Figure 15-18. Memory Activity Report

PLOT- 8

ELIGIBLES FOR SUBDISPATCH (DURATION)
 -- MIN /- AVG +- MAX



CURVE MINIMUM MAXIMUM
 MTN 7.2000E-02 4.1700E-01
 AVG 2.5007E-01 4.8654E-01
 MAX 2.1680E-00 7.9746E-03

Figure 15-19. Eligibles For Subdispatch Plot

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 PRI 82-05-21. INITIAL ENC TAPE #1469

IN SUBDISPATCH (DURATION)

— = MIN / = AVG + = MAX

[illegible]


CURVE	MINIMUM	MAXIMUM
MIN	5.900E-02	4.085E 00
AVG	3.249E 00	1.0896E 02
MAX	4.154E 00	9.2997E 03

Figure 15-20. In Subdispatch Plot

PIOT-10

ELIGIBLE FOR AND IN SUBDISPATCH (DURATION)

-- MTN /- AVG +- MAX



CURVE
MIN 7.2000E-02
AVG 2.3760E 00
MAX 6.0190E 00

Figure 15-21. Eligibles For and In Subdispatch Plot





Figure 15-22. User Swap Rate Plot

PLOT-12

USER SWAPS (SWAP AMOUNT)

-- AVG /-- TOT



CURVE MINIMUM MAXIMUM
AVG 1.0000E 00 2.0667E 01
TOT 1.0000E 00 1.1770E 03

Figure 15-23. User Swap Amount Plot

PLOT-13

USER SWAPS (DURATION)




Figure 15-24. User Swap Duration Plot

USER I/OS (DURATION)

-- MIN /- AVG +- MAX

USERS	TIME	DURATION IN INTERVAL (HS)				LINE	I/O TIMES (HUNDREDTHS SEC)					
		<1	<2	<4	<8		<16	<32	>32	>1.0E 01	>1.0E 02	>1.0E 04
0	0818:54.01	DATA NOT COLLECTED IN THIS INTERVAL.										
0	0819:14.01	1	3	4	3		<CV					
4	0819:34.01	1	0	4	3		<CV					
4	0819:54.01	0	0	2	2		<CV					
4	0820:14.01	0	0	0	0		<CV					
4	0820:34.01	0	0	0	0		<CV					
4	0820:54.01	0	0	0	0		<CV					
4	0821:14.01	0	0	0	0		<CV					
4	0821:34.01	0	0	0	0		<CV					
4	0821:54.01	0	0	0	0		<CV					
4	0822:14.01	0	0	0	0		<CV					
4	0822:34.01	0	0	0	0		<CV					
4	0822:54.01	0	0	0	0		<CV					
4	0823:14.01	0	0	0	0		<CV					
4	0823:34.01	0	0	0	0		<CV					
4	0823:54.01	0	0	0	0		<CV					
4	0824:14.01	0	0	0	0		<CV					
4	0824:34.01	0	0	0	0		<CV					
4	0824:54.01	0	0	0	0		<CV					
4	0825:14.01	0	0	0	0		<CV					
4	0825:34.01	0	0	0	0		<CV					
4	0825:54.01	0	0	0	0		<CV					
4	0826:14.01	0	0	0	0		<CV					
4	0826:34.01	0	0	0	0		<CV					
4	0826:54.01	0	0	0	0		<CV					
4	0827:14.01	0	0	0	0		<CV					
4	0827:34.01	0	0	0	0		<CV					
4	0827:54.01	0	0	0	0		<CV					
4	0828:14.01	0	0	0	0		<CV					
4	0828:34.01	0	0	0	0		<CV					
4	0828:54.01	0	0	0	0		<CV					
4	0829:14.01	0	0	0	0		<CV					
4	0829:34.01	0	0	0	0		<CV					

CURVE	MINIMUM	MAXIMUM
MIN	1.600E-02	6.512E-00
AVG	1.065E-00	2.317E-01
MAX	1.301E-00	2.062E-02

Figure 15-25. User I/Os Duration Plot

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL CMC TAPE #1469

IN SMC WAIT (DURATION)

- = MTN / = AVG + = MAX

[illegible]

CURVE	MINIMUM	MAXIMUM
MIN	2.9210E-01	1.4286E 02
AVG	2.9440E-01	1.4286E 02
MAX	2.9440E-01	1.4286E 02

Figure 15-26. In System Master Catalog (SMC) Wait Plot




Figure 15-27. Intertrace Time (Adjusted For GEAKES) Histogram

The reports that the Formatted Dump pass produces are:

- o Formatted Dump (figure 15-28) (no ID #)
- o Intertrace Time (Histogram) (figure 15-29) (ID #9).

15.5.1 TEARS Response Pass Reports

15.5.1.1 TSS Reduction Event Log. The report is a chronological listing produced as the tape is processed (figure 15-2). It indicates control changes of the Response routine, actions of the TSS Executive, and major events for individual users. For developmental use, the report also shows reduction anomalies; they are not of concern to the user. The report gives an overview of the physical makeup of the data collection period.

Each event record is formatted on one line with a time-stamp at the leftmost margin. Next to the time-stamp, for events which are related to the processing of a GMF trace type 74, the number of the trace (counting from the start of the tape) is recorded. These values can be used to locate an event on other reports or formulate timeframes for subsequent reductions. The third field is the Trace 74 subtype number. The fourth is the two-character line ID of the terminal or the four-character DRUN ID (without the suffix letter "D"). (At logon, DRUN identification is not available; the field is left blank). The next field is the name of the subsystem in control of the user at the time of the event. Following this information is a textual explanation of the event.

Examples of items on the report are:

Response control messages:

- Start of reduction tape
- Start of timeframe
- End of timeframe
- End of reduction tape.

TSS executive actions:

- Grow/reduce TSS core size
- Grow/reduce UST area
- Monitor TRACE ON/OFF.

Major user events:

- Logon/logoff
- Abnormal disconnect
- Connect command
- JDAC command.

Reduction anomalies (for developmental use):

- Program stack synchronization problems
- I/O conflicts.

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82
 DISTRIBUTION COLLECTED ON SYSTEM DSCN AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469




Figure 15-29. Intertrace Time Histogram

15.5.1.2 Response Time Logplots. The response time logplots show delay time that users experienced waiting for a response from TSS. They are:

- o Response Time For All Users (figure 15-3)
- o Response Time For Users Not Requesting More Core (figure 15-5)
- o Response Time For Users With Core Request During Line Idle (figure 15-6).

(Logplots are described in more detail in section 15.3).

The All-Users logplot shows, from time interval to time interval, the minimum, average, and maximum amount of time a user waited for a response. The user who waited longest is identified under the column LINE. An individual wait may extend over more than one interval. In this case the time value used on a plot line will be the total amount of time since the start of the line-idle state (not just the portion of time in that plot line period). This logplot should indicate periods of time when users were dissatisfied with TSS performance.

The information used on the All-User logplot is also used on two other logplots. These logplots help to indicate times when the source of response time degradation is due to swap area competition. The wait times are distributed on the second and third response time logplots on the basis of a request for core during the line-idle period (e.g., call to subsystem, or swap in). Differences in these plots (high response times on figure 15-6 and lower response times on figure 15-5) would be explained by Swap Area Competition. The three figures in this document would indicate that swap areas were not a problem. In fact, the only terminal experiencing difficulties appears to be terminal ID whose response time keeps on increasing during each interval.

15.5.1.3 Subdispatch Time Logplots. The subdispatch time logplots show delay time for subsystems in obtaining service by GCOS and TSS Executive. They are:

- o Total Time in Subdispatch Queue (figure 15-7)
- o Time in Subdispatch Queue Waiting Service (figure 15-8)
- o Processor Time in Subdispatch (figure 15-9).

The left-hand side of the three reports are the same from the "USERS" column to the "I/O-BSY" column. The "USERS" column is the count of users in the time-interval ending at "TIME". During each interval, the total number of subdispatches performed is contained in the column "#SD". The value in the "CPU" column is the percentage of total amount of time in subdispatch queue during which the processor was executing TSS user subsystem code (the remainder being queue overhead time). The percentage of these subdispatches which ended as a timer runout fault (vs. a DRL fault) are indicated in the column "TROUTS". The percentage of subdispatches executed for subsystems which, concurrently, had an outstanding I/O to the terminal is reported in

the "I/O-BSY" column. The "LINE" field indicates the line-ID of the user who contributed the datum for that line of that plot with the greatest magnitude (see below for an explanation of the data that is plotted on each of the reports).

The right-hand side of the "Total Time" logplot shows, from time interval to time interval, the minimum, average, and maximum duration of all subsystem subdispatches. The time value used as data for this logplot is measured from the event at which the TSS executive creates a "ready entry" in the subdispatch table, to the event at which the executive recognizes that GCOS has completed processing by changing the entry to a "fault entry".

The two constituents of each time value contributing to the "Total Time" logplot are segregated to produce the right-hand sides of the "Waiting Service" and "Processor Time" logplots. The first of the two values is the sum of the time that the "ready entry" waited for allocation by the CPU allocator, and the time that the "fault entry" waited to be detected by the TSS executive. The second value is the actual processor time used by the subsystem (between "ready entry" and "fault entry" states).

15.5.1.4 CPU Times Allocated to TSI-Exec and TSI-Subdispatch Report. The CPU time allocation chart (figure 15-4) shows the distribution of processor time attributable to the TSI Executive, to the TSS subdispatch mechanism, to all non-TSS customers, and to the idle state. This report is produced only if the CPU Monitor (CPUM) was in execution when the data tape was written. Data collection for each print line is independent of other TEARS procedures, however the trigger to print a line is synchronized with the response time logplots. Thus, the accumulation of CPU data into this report will "get ahead" of other TEARS reports when a large gap exists between TSS traces. Here's why:

Time driven reports, including this one, are based on TSSM trace times alone, not on CPUM trace times. If a large gap exists between TSSM traces, one or more successive "times-to-print" may expire without any printing. During the TSSM trace gap, however, CPUM traces continue to be generated and accumulated into this report. When TEARS eventually finds the long delayed TSSM trace time stamp, it triggers plot lines at successive plot intervals until it catches up with the time stamp. At the first such trigger, all the accumulated CPU data is included in the print line; successive triggers will, of course, find no additional CPU data to print. The mechanism just described explains (1) the blank lines in the CPU Time Allocation report, and (2) the greater-than-expected CPU times immediately preceding a blank line. The mechanism was deliberate so that gaps in TSS service are obvious.

Because this mechanism differs from the other TEARS time-driven reports, and to fully describe the scope of the data for each print line, the first three columns of the report are devoted to time data. In Figure 15-4, our sample report, the first item shows the expected "time-to-print" (PLOT TIME

REFERENCE). The next two columns show the time of the last CPU trace processed in the previous print line (FIRST T70 TIME) and the last CPU trace processed in the current print line (LAST T70 TIME) -- that interval, first to last, expresses the actual elapsed time over which the CPU data pertains. The next item is the number of CPU traces on which the data is accumulated (NUMBER OF T70S). The right half of the report is a tabular chart showing the CPU time, in seconds, attributed to the TSS Executive (TSS-EXEC), to the TSS subdispatch mechanism (TSS-SDSP), to the idle state (IDLE-TIME), and to all non-TSS consumers (OTHER-CPU). The last numeric column is a total (TOTAL) of the preceeding four columns. The graphic to the extreme right shows the deviation of the total CPU time from that expected based on the plot interval. Each "+" to the right of the axis ("I"), shows a 10% deviation "greater than expected". An asterisk replaces the "+" when the deviation is 100% or greater. Deviations in the other direction are shown as "-" symbols to the left of the axis.

For a normally functioning system under moderate load, the number of CPU traces approximates 1 for each 2 CPU-seconds. If the plot interval is so small or CPU activity so light that fewer than about five CPU traces are collected for each print line, the report will show false abnormality. Thus, an "abnormal" report may result from (1) low CPU activity, (2) too short a plot interval, or (3) gaps in TSS service.

15.5.1.5 TSS Subtraces Encountered, and Derails Executed by Users. The two summary histograms "TSS Subtraces Encountered" (figure 15-10) and "Derails Executed by Users" (figure 15-11) are incorporated in the pass to show the presence and volume of those traces/events within a timeframe.

15.5.2 TEARS Emulation Reports. The TEARS Emulation phase is both trace driven and elapsed-time driven; the distinction also separates two categories of reports for the user. Reports driven by trace events correspond with mimicking TSS - they give evidence of user difficulties and delays or of TSS Executive difficulties and delays. The reduction program remembers, on a short term basis, what each TSS user was attempting to do, the obstacles encountered in the attempt, how long the attempt was ongoing, and some of the computer resources used in the attempt. Report entries are made when some threshold, usually based on time delay, is exceeded or when a service denial, failure, or error thwarts the user. The program also measures the frequency with which the TSS Executive did its work - a report entry follows a delay threshold.

Reports driven by elapsed time measure several TSS functions without regard to any particular user -- they report measures of the subdispatch mechanism, memory activity, and the user disk I/O functions among others.

Several reports participate in a superior-subordinate arrangement; only selected entries in the superior report cause an entry in the subordinate report. The Users' Map is subordinate to the Exception Message Report and the TSS Core Map is subordinate to the TSS Delay-User Delay Report. They

are event driven reports. One elapsed-time driven report, the Memory Activity report, is subordinate to the logplot of Intertrace Gap but only for synchronization convenience.

15.5.2.1 Exception Message Report. When a request for some system service, TSS or GCOS, is denied or terminates with an error, an entry is made in this Exception Message Report (figure 15-12). This report also controls (triggers) item entries in the Users' Map, section 15.5.2.2, for several of the errors/denials. The entries are categorized by function/origin into four classes: TSS, FMS, Derails, and the inevitable Miscellaneous.

An additional entry identifies those users who consume excessive CPU resources (no denigration of the user is intended, nor should you infer it). The intent of the entry is (1) to correlate an apparently excessive response time pattern with a user engaged in CPU intensive work (with occasional messages to his terminal), and (2) to identify the magnitude of such work.

A contrived example of the Exception Message Report is given as figure 15-12. It provides one example of each possible error message. Reading from the left, you can see the time (TOD) and sequence number (RECNO) of the trace which captured the denial/error. Time of day is represented as HHMM:SS.sss. The current trace sequence number is relative to the first TSS trace following GMC start. This is followed by the message proper, shown under the heading EXCEPTION MESSAGE. We list the user's terminal ID and his USERID, if applicable, followed by the error message. Error messages are taken directly from Honeywell publications, and some are cryptic. For excessive CPU usage, an exception is declared when the accumulated CPU time, since swapping in, exceeds 64,000 clock pulses; the CPU time is checked each time a user is removed from the subdispatch fault queue.

To reduce repetitious entries for the same user, such entries are inhibited so as to print not more than once per second. The graphic to the right of an entry shows one dot, up to 30, for each inhibited "print" of the previous entry for that user. An asterisk following the time of day indicates that an accompanying entry in the Users' Map report has been made.

15.5.2.2 Users' Map. This report (figure 15-13) is driven by the Exception Message Report for several exception entries. Recall that in the parent report (section 15.5.2.1), the triggering entry is marked with an asterisk following the time of day. The same entry from the parent report is again printed on the Users' Map as the reason for the map. Following the reason is the users' map proper. It displays, in inverse core location order, all current TSS users and their status. We show the user sequence number (USR), his full address (USTOCT) in octal, his line/terminal identification (LID) and his user identification (USER-ID). The user sequence number is a TEARS-generated subscript into a set of program tables, while LID shows the line ID written either as id| for terminal sessions, or as NNNN for DRUN sessions (where NNNN is the DRUN number). Following the user identification

is user state data -- how long the user has dwelled in his current state (TIME-IN) in milliseconds and his current state (STATE) in mnemonic form (viz, table 15-1). If terminal I/O is in progress, a "T" (for true) is printed in the next column (TRMIO). We then show the number and mnemonic of the last recorded trace (LAST-TRACE) and the last recorded derail (DERAIL) for the user. The following items report the last recorded program size (SIZE) and an indicator that core is currently assigned ("T" for true). The last item shows the entire program stack (PROGRAM-STACK); an asterisk following the top of stack denotes a user in build mode.

15.5.2.3 TSS Delay - User Delay Report. This is the most complex TEARS report (figure 15-14). It is a time-ordered merge of two separate event histories so that user delays rooted in TSS delays become obvious. It also reveals some TSS delays caused by user service requests, the most abundant being security package requirements (.FSO49 during DRL DEFIL). This report also controls (triggers) item entries in the TSS Core Map, section 15.5.2.4, for users waiting core too long.

From the left in figure 15-14, the columns TOD, RCNT, and CUR are common to both TSS delay and user delay entries; they show the current time of day, the sequence number of the current TSS trace, and the subtype of the current TSS trace. Time of day is represented as HHMM:SS.sss. The current trace sequence number is relative to the first TSS trace following GMC start, and the subtype tells which of the 105 traps caused the trace. An asterisk following the current trace number identifies a trace performed in courtesy call. The remainder of the report describes a TSS delay, a user delay, or a concurrency of both types of delay. In figure 15-14, examples of all three-type entries appear. To highlight specific entries in figure 15-14 for further discussion, those entries are enclosed in braces and marked with a key letter (A, B, ...). Reference to these specific entries is made by a phrase such as "Key Z...".

Key A is an example of TSS delay reporting - it shows a 1-second delay (DLAY) between trace subtypes 18 and 61 (PRV, CUR), where the delay is printed as rounded seconds. Subtype 18 is a "Gewake with subdispatch busy" and subtype 61 is an "Enter processor allocation." TEARS program logic identified the subtype 18 trace as the "cause" by executive delay (XEC); this means that the TS1 executive gave up the processor with a GEWAKE and was delayed at least 1 second longer than the GEWAKE time. Thus the delay was external to TSS, and cannot be further localized. The indicators of external-to-TSS delays are the following subtypes (PRV):

- 15, "Gewake - no users"
- 17, "Gewake until subdispatch done"
- 18, "Gewake with subdispatch busy"
- 72, "Issue GEMORE for memory"
- 102, "Issue remote inquiry GEROUT".

and (CUR):

- 61, "Enter processor allocation"
- 89, "Enter routine to process queue".

The central portion of the report entry (columns DRL through USR-I/O) is a tableau that links functional areas with last known trace subtypes for any user. The functional areas are the column headers and represent derail service (DRL), File Management Supervisor service (FMS), and disk service (USR-I/O). The functional areas are ones in which TSS delays have been observed or are possible; column entries are the trace subtypes associated with such a TSS delay. If a linking is made, it is only a suggestive or circumstantial cause of delay. Several disparate linkings are possible, as is a failure to make any linking at all. The possible linkings are as follows:

For derails (DRL), the unique entry is trace subtype 23, "Process derail."

File Management Supervisor delays (FMS) are associated with two trace subtypes - subtype 27, "Issue MME GEFSYE" and subtype 45, "Call .MFS19 for file grow."

Users requesting disk I/O (USR-I/O) have caused TSS delays with trace subtype 24, "Request file I/O", and subtype 41, "Load subsystem with DRL RESTOR."

The right half of the report (Key A remains the referent) is a snapshot of all current users. The items reported for each user are "as of" just after the current trace was processed. Each user is identified by user sequence number (UST), terminal identification (LINE), and user identification (USERID). Following this, the time the user has dwelled in his current state (DLAY) is shown as rounded seconds and the current state (STATE) is printed. Following the state, an asterisk in column T marks a user engaged in terminal I/O. The next two columns display the last recorded trace subtype (PRV) and the last recorded derail type (DERAIL) for the user. Next, the last known program size (SIZE) is dimensioned as words-of-memory/1024. The last item shows the top three (of a possible five) members of the user's program stack (PROGRAM STACK). An asterisk following the top of stack denotes a user in build mode.

Notes:

- (1) In this example, Key A, the total time between the listed traces was 1 second (rounded); program logic detected the delay as excessive because the time between listed traces was at least 1 second (not rounded) longer than the GEWAKE time.
- (2) In single processor systems, the delay would not have been detected, since GEWAKE time is not available. (Program logic would have taken the GEWAKE time to be very large. The logic could as well have taken the GEWAKE time to be very small. We chose the former alternative to eliminate false alarms).

(3) It is possible that TSS traces performed in courtesy call appeared within the interval bracketed by "PRV,CUR"; they are ignored within (the TSS delay part of) this report.

Key B is an example of user delay reporting. This type entry is made only for users who are delayed in a vulnerable state (refer to table 15-1). Further, an entry is made for a user only when a trace attributable to that user is the current trace. Except for the state "waiting for memory" (mnemonic WTMEM), a delay is detected when state occupancy exceeds 1 second. For the state "waiting for memory", a delay is detected when state occupancy is 1 second longer than the wait time computed on program size.* The items reported for a delayed user are "as of" just before the current trace was processed. The data items are the same as described above for the user snapshot portion of TSS delay entry.

Key C is an example of user delay reporting while waiting memory. The asterisk following the time-of-day indicates that an accompanying entry in the TSS Core Map Report has been made.

Key D is an example of a concurrent TSS delay and user delay where TEARS logic failed to identify the cause (.FSO49 in DRL DEFIL) but inspection can. The trace subtype which caused the problem was a "derail" (trace subtype 23) and the derail was a DEFIL. After the DRL DEFIL began, a "terminal I/O complete" (subtype 87) in courtesy call was trapped, and that trap caused TEARS logic to "forget about" the derail although the state was preserved. The clues to the interpretation are the values under user state, DRLSR, the last known derail, DEFIL, and the fact that the user delay and TSS delay are the same. (Logic error may be corrected in delivered version.)

Key E shows two user delay entries for the same user in the same state, but with a user option in force. To preclude repetitious entries for a user in the same state, a user option (option(onchange)) was selected here. It inhibited all but the first print following delay detection and the last print just prior to state change. The plus sign following the state indicates (a) the user's state will change as a result of the current trace and (b) the delay in this state for this user has been printed once before (but for a lesser delay time).

15.5.2.4 TSS Core Map. This report is driven by the TSS Delay - User Delay Report logic when it detects a user waiting memory too long. Recall that in the parent report (section 15.5.2.3) the triggering entry is marked with an asterisk following the time-of-day. The same entry from the parent report is again printed on the TSS Core Map as the reason for the map. Following the reason is the core map proper. It displays, in core location order, the

* TS1 computes an alarm time as a function of each user's program size. It will take no extraordinary action to obtain memory for a user until the time has expired. TEARS uses the same function (plus 1 second) to detect a delay when waiting for memory.

allocation of all the user reserved swap area. The allocation printout is condensed so that each line describes one user in core along with any free memory contiguously above him.

Reading from the left in figure 15-15, find the user sequence number (USR), and the starting address of his swap area (BASE) relative to TS1. The amount of core assigned (SIZE) is followed by the unused memory (FREE) contiguously above it. If no user occupies the lowest core slot, the first print line shows that as a HOLE in core. The remaining items of the print line describe the user and what he is doing.

The user is identified by his line identification (LID) and user identification (USER-ID). The next three items are timers dimensioned in milliseconds: the user's time in current state (TSTAT); time since last completed "swapping in" to core (TCORE); and accumulated CPU time (TCPU) since last completed swapping in to core. The next item is a measure of disk resource use -- since disk connect time is not available from within TSSM, we collect and print instead the number of disk I/Os (DIOS). This is followed by the mnemonic for the last recorded derail (LSTDRL), the entire program stack (PROGRAM STACK), the complete address of the user's UST (USTOCT) in octal, and the user's current state (STATE). If terminal I/O is in transmission, a "T" (for true) is printed in the next column (TRMIO). The last item is the mnemonic of the last recorded trace (LAST TRACE) for the user.

Notes:

(1) Core allocated to a user as Extended Buffer Memory is indicated as such by printing "---ebm---" in place of the program stack (PROGRAM STACK).

(2) The items TCORE, TCPU, and DIOS for a user just swapping in may be misleading -- because a user swapping in has core assigned as shown in the core map, but he is not considered in core. The three items mentioned are "as of" the completion of his previous swap into core.

The Core Map must be interpreted in conjunction with its parent to avoid wrong inference. In particular, a long delay waiting memory may be rooted in a TSS delay and thus is not a sign of memory allocation trouble. Also, not every excessive wait for memory (shown in the parent) results in a core map entry; to avoid redundancy, the core map print is completely inhibited if TEARS logic shows no change of core allocation since the last printed map. Further, the option described in section 15.5.2.3 that limits printing to "first and last occurrence" also limits the core map.

15.5.2.5 Error Message Report. The Error Message Report identifies conflicts between how the TSS works and how the TEARS designers expected it to work. Errors reported here are not fatal to the reduction pass and will not invalidate analysis. Nevertheless, the errors should be reported to

C751 for program correction. The report format is identical with the Exception Message Report described in subsection 15.5.1.1. Figure 15-16 is both a sample report and a list of the known errors as of the version number and compilation date shown.

15.5.2.6 Intertrace Gap - Duration. This logplot is an attempt to graph the time between TSSM traces while subtracting out the expected duration of any TSS Gewakes. The goal is a graph in which large gaps represent periods of delayed service and not periods in which TSS was idling by design. To avoid contamination of the data by traces performed in courtesy call during the periods TSS was idling in Gewake, all courtesy call traces are excluded. Figure 15-17 is an example of this report. It plots the minimum, average and maximum time between traces that are not performed in courtesy call. Those gaps which include a Gewake are corrected by subtracting out the expected Gewake time. The plot also shows the number of users, the time of the plot line, and seven counts indicating the number of gaps in seven ranges of time. For example, the 4 column lists the number of gaps that were less than 4 milliseconds; the 8 gap, between 4 and 8 milliseconds, and so on. The same data set used here is also source for the histogram Intertrace Time described in subsection 15.5.2.16.

For single processor systems, the expected Gewake time is not available - the gap time following a Gewake is taken as zero.

15.5.2.7 Memory Activity Report. The TSSM subtraces 63 and 65 through 75 form the bulk of exceptional TSS Executive service in support of memory requirements for the user. Figure 15-18 is an example of this report. It is a time-driven tabular report; it is printed in synchrony with the logplot of intertrace times for convenience alone. Each print line shows the count of each of the traces 63 and 65 through 75 that were trapped during the plot interval. The traces are:

- 63, "Enter swap decision processing"
- 65, "Consider TSS size increase"
- 66, "Initiate size increase for urgent user"
- 67, "Set up fence for urgent user"
- 68, "Set up urgent user class memory reserve"
- 69, "Force swap"
- 70, "Terminate force swap"
- 71, "UST area increase by 1K"
- 72, "Issue Gemore for memory"
- 73, "Gemore successful"
- 74, "Memory release"
- 75, "Gemore refused or reduction not possible".

Together with the two logplots described in sections 15.5.2.11 and 15.5.2.12, the Memory Activity report provides a provisional and incomplete picture of TSS memory allocation service.

15.5.2.8 Eligibles for Subdispatch - Duration. This report plots the minimum, average and maximum wait-subdispatch-time (WTSUB state) in milliseconds for all users. The WTSUB state begins when TEARS logic determines that the user is eligible for subdispatch; it ends when the user is placed in the subdispatch queue. Each line of the logplot is a window of a length determined by user option (default is 20 seconds). The logplot shows the number of users, the time of the plot line, and the line ID of the user with the maximum time. It also shows seven counts indicating the number of waits in seven ranges of time. For example, the 1 column indicates the number of waits that were less than 1 millisecond. The 2 column indicates the number of waits that were between 1 and 2 milliseconds, etc. Figure 15-19 is an example of this report.

15.5.2.9 In Subdispatch - Duration. This report plots the minimum, average and maximum subdispatch time (SUBDS state) in milliseconds for all users. The SUBDS state begins when the user is placed in the subdispatch queue; it ends when TS1 finds the user in the fault queue. Each line of the logplot is a window of a length determined by user option (default is 20 seconds). The logplot shows the number of users, the time of the plot line, and the line ID of the user with the maximum subdispatch time. It also shows seven counts indicating the number of subdispatches in seven ranges of time. Figure 15-20 is an example of this report.

15.5.2.10 Eligible for and in Subdispatch - Duration. This report plots the minimum, average and maximum combined time of wait for subdispatch (WTSUB state) and in subdispatch (SUBDS state) for all users. An example of this report is given in figure 15-21. This logplot defines the length of a combination event in three ways. When a wait-subdispatch is followed by a subdispatch, the length of the event is from the start of the wait to the end of the subdispatch. If a wait-subdispatch is followed by something other than a subdispatch, the length of the event is from the start of the wait to the end of the wait. If a subdispatch is preceded by something other than a wait-subdispatch, the length of the event is from the start of the subdispatch to the end of the subdispatch. Each plot line is identical in format to both the Eligibles for Subdispatch, and the In Subdispatch logplots.

15.5.2.11 User Swaps - Swap Rate. This report plots the rate of core swapped, in K-words per second, for all users. Core to core swaps are not included. Each plot line is a window of a length of time determined by user option (default is 20 seconds). The logplot shows the number of users and the time of the plot line. The rate is defined as the total swapped core divided by the window interval. This logplot is not keyed on users' states but on swap events. Figure 15-22 is an example of this report.

15.5.2.12 User Swaps - Swap Amount. This report plots the average and total core swapped, in K, for all users. Each plot line shows the number of users, the plot line time, the line ID of the user with the maximum core swapped, and seven counts indicating the number of swaps in seven block

sizes of core. This logplot is not keyed on user states but on swap events. Figure 15-23 is an example of this report.

15.5.2.13 User Swaps - Duration. This report plots the minimum, average and maximum swapping time in hundredths of a second and does not differentiate between swap in (SWPIN) states and swap out (SWPOU) states. The state begins when TSI initiates the swap; it ends when the swap is complete. Along with the number of users, the time of the plot line, and the line ID of the user with the maximum swapped time, this logplot shows seven counts indicating the number of swaps in seven ranges of time. Figure 15-24 is an example of this report.

15.5.2.14 User I/Os - Duration. This report plots the minimum, average and maximum time all users were in disk I/O (DSKIO state) in hundredths of a second. Each plot line shows the number of users, the plot line time, the line ID of the user with the maximum I/O time and seven counts indicating the number of I/Os in seven ranges of time. Figure 15-25 is an example of this report.

15.5.2.15 In System Master Catalog Wait - Duration. This report plots the minimum, average and maximum time all users were in SMC wait (WTSMC state) in hundredths of a second. Each plot line shows the number of users, the plot line time, the line ID of the user with the maximum SMC wait time and seven counts indicating the number of SMC waits in seven ranges of time. Figure 15-26 is an example of this report.

15.5.2.16 Intertrace Time (Adjusted for TSS GEWAKES). This report is a histogram of the time between TSSM traces that are not performed in courtesy call. The intertrace time is corrected for Gewakes as explained in section 15.5.2.6. Data for the report is presented in rounded milliseconds. A sample report is given as Figure 15-27.

15.5.3 Formatted Dump Reports

15.5.3.1 Formatted Dump. The formatted dump is a debug or investigative tool for the technical user. It decodes most of the data fields which comprise each trace, correlates UST-bearing traces with the corresponding line ID and prints a one-line expansion of the trace. Were the report unlimited, one GMC data tape would generate enormous output. To filter this mass of data, the user options TIMEFRAME, TRACES, and LINES may be used as described in subsection 15.6.

Since the data content of the 105 trace 74 subtypes varies so widely, a generalized column descriptor heads the report. From left to right, each trace expansion shows: the current time-of-day (TOD) as HHMM:SS.sss; the sequence number of the current TSS trace (REC-CNT) beginning with the first TSS trace after GMC start; the name or mnemonic and subtype number (MNEMONIC & #) identifying the current TSS trace. If the trace bears UST information, the two columns, UT # and LINE, uniquely identify the user to which the

trace is attributable -- the UST # is an internal subscript into a set of program tables, while LINE shows the line ID written either as id| for terminal sessions, or as NNNN for DRUN sessions (where NNNN is the DRUN number). The columns headed LITERAL and NOTES may contain decoded interpretations of various data fields within the trace -- in general, 12-character numerals and numerals with leading zeros are octal except that the data item lflg2 (i.e., .LFLG2) is binary. Any trace performed in courtesy call shows an asterisk following the current time-of-day. The subtype 90 trace is not decoded directly -- after a "90 sequence" does occur (and also at the start of a formatted dump), a list of current users is shown in line-item format:

current-time LINE id IS "user id" --LOGON AT logon-time

The subtype 105 trace contains too much data for a 1-line expansion. Currently, it is only cursorily expanded - LITERAL shows the trace A-register.

15.5.3.2 Intertrace Time. This report is a histogram of the time between TSSM traces. Data for the report is presented in rounded milliseconds. The report is of identical format to that described in section 15.5.2.16; a sample report is not repeated here.

15.6 Reduction Run-Time Options

The general card format for an option request is as follows:

COMMAND(PARAMETER-LIST)

where COMMAND is a keyword specifying what action is to be taken and PARAMETER-LIST, if required, provides data to accomplish the action. The COMMAND is interpreted through its sixth character, so the input must match the first six characters of a valid command. Shorter commands must match exactly.

Fourteen commands are recognized by the card input interpreter. The list below shows each command, the action to be taken, and in parentheses, the default implications:

- o CPRINT Print all time driven reports at a specified repetition period. (Print at 60-second intervals for response pass reports; 20 seconds for emulation pass reports).
- o DEBUG Activate the debug statements in one or more modules. (All debug statements are inactive).
- o DUMP Reserved for development.

- o HISTOGRAM Modify a histogram report. (Histogram report parameters are as shown in table 15-2).
- o LINES Restrict reduction to include or exclude specified line/DRUN ids. (No restrictions).
- o NEW Prepare to accept another, or repeat, set of GMC data tapes and another set of option alteration commands for a following reduction run. (Reduction is complete when the current set is reduced).
- o ON (1) Activate the reduction routines for a specified pass over the GMC data tapes. (Response time pass is active; emulation and formatted dump passes are inactive).

(2) Turn on one or more specified reports. (All reports are on).
- o OFF Turn on one or more specified reports. (All reports are on).
- o OPTION Enforce a special option. (All special options are inactive).
- o PLOT Modify a logplot. (Plot parameters are as shown in table 15-3).
- o STOP Stop reduction after processing a specified number of physical tape records. (No limit).
- o TIMEFRAME (1) Accept one to five time windows for overall data reduction or for specified reports. (Data reduction and report output are not time limited).

(2) Accept one to five time windows for debug output. (Debug output is not time limited).
- o TITLE Accept a new title to be printed on all reports. ("TSS REDUCTION, VERSION 7.2 - 1.000, 1 JULY 82" is printed as the banner).
- o TRACES Restrict reduction to include or exclude specified TSSM subtrace types. (No restriction).

15.6.1 Parameter List. The parameter list is a sequence of data items separated one from the next by a field of spaces. Any one space in the field may optionally be replaced by a comma or colon. A data item may be an alphanumeric, a quoted string, an integer, a real number or a null. Strings

Table 15-2. Histogram Default Parameters

<u>REPORT ID</u>	<u>TRAILER FLAG</u>	<u># INTERVALS</u>	<u>LOW VALUE</u>	<u>INTERVAL SIZE</u>
1	OFF	54	1	1
2	OFF	54	54	1
7	OFF	39	-10	1
8	OFF	40	28	1
9	ON	54	0	64

Table 15-3. Plot Default Parameters

<u>REPORT ID</u>	<u># POINTS TO PLOT</u>	<u>Y-MINIMUM</u>	<u>ORDERS-MAGNITUDE</u>
4	UNLIMITED	1.0	3.0
5	UNLIMITED	1.0	3.0
6	UNLIMITED	1.0	3.0
10	UNLIMITED	0.1	6.0
11	UNLIMITED	0.1	6.0
12	UNLIMITED	0.1	3.0
15	UNLIMITED	0.1	5.0
14	UNLIMITED	0.1	4.0
15	UNLIMITED	1.0	4.0
16	UNLIMITED	0.1	5.0
17	UNLIMITED	1.0	3.0
18	UNLIMITED	1.0	4.0
19	UNLIMITED	0.1	4.0
20	UNLIMITED	0.1	4.0
21	UNLIMITED	1.0	5.0

are enclosed in single or double quotes. Real numbers include the standard FORTRAN forms: for example, 12.34, 1234E-2, 1.234+1 are all accepted as the same number. A null is two consecutive commas or colons; it stands for a missing data item and for the separator fields on either side of the missing item.

The data acronym RID stands for a one- or two-digit integer specifying the report identification number of the report to be affected. Report identification numbers are shown in section 15.5. The specifications "integer-W", "alphanumeric-X", "real-Y", and "string-Z" require user input of an integer number in a field at most W wide, an alphanumeric in a field at most X wide, a real number in a field at most Y wide, and a quoted string in a field at most Z wide (not including the quote delimiters), respectively. Alphanumerics and strings exceeding the specified field width are accepted but are truncated to the right. Integers and reals exceeding the specified field width are rejected as illegal data. Integers may not contain a decimal point. Reals must contain a decimal point or an exponent. Reals are limited, for all data input, to a maximum of 14 characters including sign, decimal point and exponent. Reals without exponent are further limited to a maximum of 11 characters, including leading sign and decimal point. Numeric data items shorter than the specified field width are acceptable in all commands, but a null item is acceptable only where provided for.

15.6.2 Command Syntax. Each command with its parameter list must be complete on a single card, but more than one command may be placed on each card if spaces separate them. The parameter list may either be contiguous with the command keyword or spaces may intervene. The syntax for each command is given below.

TEARS permits reduction of the GMC data tapes in one of three phases: response time analysis, emulation analysis, or formatted dump analysis. The response time analysis routines are on, and the other phase routines are off, by default. The following commands may be used to exercise control over the reduction:

- o DEBUG Activate all debug; if emulation phase, load special debug link.
- o DEBUG(MODULE-NAME-1,...,MODULE-NAME-N)
 Activates the debug statements in the explicitly named modules, where MODULE-NAME-i, alphanumeric-6, is the FORTRAN name given to the subroutine, function or link. The modules currently nameable and the intent of the debug action, in parentheses, are shown below.
 - | BUGDMP - additional debug link for emulation phase.
(Performs an interlineated formatted dump and finite-state-machine dump of the 1000 traces preceeding a fatal failure of the emulation phase logic).

| CHANGER - decodes trace data for formatted dump. (Allows comparison of TEARS decoding with trace data).
 | CHKPOINT - saves state of the emulation phase finite-state-machine to random file. (Verifies action of routine).
 | DOTRAON - intermediate control routine for formatted dump. (Shows correlation of user sequence number (a major TEARS data structure) with TSS UST address).
 | DOTRALIMIT - not a module but a control. (Limits formatted dump to one sample of each TSSM trace subtype).
 | GETOKE - lexical analysis of card input. (Verifies action of routine).
 | LOGPLOT - formats and prints logarithmic bar chart. (Prints arithmetic values in lieu of bar chart).
 | NXTRECRD - tape input handler. (Verifies action of routine).
 | SAMSON - controls emulation phase when within a reduction timeframe. (Performs formatted dump during emulation phase).
 | SETSTATE - the finite state machine of the emulation phase. (Provides additional data in case of non-fatal errors).
 | STATECHART - not a module but a control in the emulation phase. (Prints one table, at each plot interval, of the state transitions of the finite state machine during the interval).
 | TIMESET - sets up array of start-stop times for user-specified reduction windows. (Verifies action of the routine).
 | T90 - initializes reduction variables in accordance with TSSM trace subtype 90. (Verifies action of the routine).

The one shot subroutine GMFEXC, which reads the initial tape record, is executed before card input. It may be debugged by setting the program switch word, bit 11.

- o HISTOGRAM(RID, TRAILER-FLAG, NUMBER-OF-INTERVALS, LOW-VALUE, INTERVAL-SIZE)

Turns on the histogram with report identification number RID, integer-2. Prints/does not print the histogram summary if TRAILER-FLAG is ON/OFF, alphanumeric-3. The new number of histogram intervals is NUMBER-OF-INTERVALS, integer-3. The new low value and interval size are LOW-VALUE and INTERVAL-SIZE, both integer-10. Note: In the current version, the number of intervals may not exceed 54.

The following two options are as above, but without change to the current specifications of those parameters not appearing in the list.

- o HISTOGRAM(RID , , NUMBER-OF-INTERVALS, LOW-VALUE, INTERVAL-SIZE)
- o HISTOGRAM(RID, TRAILER-FLAG)

- o `LINES(SENSE-OF-RESTRICTION, "ID-1", ..., "ID-N")`
Excludes from, or includes in, the formatted dump which only those lines/DRUNs listed. Exclusion is obtained when the SENSE-OF-RESTRICTION is DONTDO, alphanumeric-6; inclusion, when the SENSE-OF-RESTRICTION is DOONLY, alphanumeric-6. Up to 10 lines or DRUNs may be listed in 1 or several commands, but if the SENSE-OF-RESTRICTION changes, previous commands are discarded. Terminals/lines are specified by a string "ID", string-2; DRUNs, by a string "NNNN", string-4. String lengths must be exact.
- o `NEW("TAPE-NUMBER")`
Flags TEARS to stop reading card input options and to proceed with data reduction. It also informs the program that another data reduction pass is to follow, and that the following pass is to lead off with the tape number specified by "TAPE-NUMBER", string-6. For that following reduction, the default options will prevail unless modified by option alteration cards following the NEW command. This command must be the last command on a card; commands following on the same card are discarded. Note: TAPE-NUMBER may optionally be of the form alphanumeric-6, NEW(AJ283) for example.
- o `OFF` Turns off all switchable reports.
- o `OFF(RID-1, ..., RID-N)`
Turns off the reports specified by report identification numbers RID-i, integer-2.
- o `OFF(CATEGORY)`
Turns off all logplots or histograms depending on CATEGORY: PLOTS or HISTOGRAMS, both alphanumeric-2.
- o `ON` ON options use the same parameters as those shown above for the OFF command, with appropriate reversal of intent.
- o `ON(PHASE-OF-REDUCTION)`
Turns on, or engages, exactly one of the three phases of reduction depending on the value of PHASE-OF-REDUCTION: RESPONSE, EMULATE, or FORMATDUMP, all alphanumeric-6. The response phase is on by default; all others are off. Only one phase may be accomplished on a single pass over the data tapes.
- o `OPTION(CHOICE)`
Modifies action of the reduction program depending on the following values of CHOICE, each alphanumeric-6:

| ALLCORE

Expands the two emulation phase plots that depict User Swap amount and rate so as to include core-to-core movement. The expansion is imperfectly implemented and the option is not recommended.

| LOGSCALE

Inserts intermediate numeric values on the y-axis for all logplots.

| MINUSDISKIO or -DISKIO

Adjusts individual response times of the response phase by subtracting the time spent accomplishing each disk I/O in the line-idle period.

| MINUSSUBDISPATCH or -SUBDISPATCH

Adjusts individual response times of the response phase by subtracting the CPU time in subdispatch expended during the line-idle period.

| NOABORT

Reserved for development.

| ONCHANGE

Limits the trace driven report "TSS Delay - User Delay" of the emulation phase to print user delay items not more than twice: first, when the delay is initially recognized, and second when the delay is about to end.

| TK4014

Reserved for development.

o PLOT(RID, NUMBER-OF-POINTS, Y-MINIMUM, ORDERS-OF-MAGNITUDE)

Turns on the logarithmic plot with report identification number RID, integer-2. The new number of points to plot is NUMBER-OF-POINTS, integer-10, maximum. The y-axis is bounded below by Y-MINIMUM, real-14, and above by $Y-MINIMUM * 10^{ORDERS-OF-MAGNITUDE}$, where ORDERS-OF-MAGNITUDE is real-14. (The exponent is truncated to integer before use; if that integer is not between 1 and 7 inclusive, it is arbitrarily set to 7.).

The following two options are as above, but without change to the value of those parameters not specified.

o PLOT(RID, NUMBER-OF-POINTS)

o PLOT(RID, , Y-MINIMUM, ORDERS-OF-MAGNITUDE)

- o STOP(RECORD-COUNT)

Stops the data reduction pass after reading RECORD-COUNT, integer-10, physical records, maximum. One fewer records will be processed.
- o TIMEFRAME(RID, START-TIME-1, STOP-TIME-1,...,START-TIME-N, STOP-TIME-N)

Turns on report with report identification number RID, integer-2. Accepts up to five start/stop pairs: START-TIME-i, STOP-TIME-i where each component of the pair is real-14. Times are input in psuedo-decimal form:
 HHMM.SSsss
 where HH is the hour in 24-hour format, MM is the minutes of the hour, SS is the seconds of the minute, and sss is the fraction of seconds. Up to five start/stop pairs may be specified for a given report. The count of five may be achieved in one or several commands, however, the lexical order must conform with the temporal order to get the results wanted. The final pair in a command may be abbreviated -- if only a start time is given, the stop time is permanently unbounded. If RID is zero, the time frame(s) bound the entire reduction, but the status of individual reports is unaffected. Histograms may not be individually delimited with the command.
- o TIMEFRAME(DEBUG, START-TIME-1, STOP-TIME-1,...,START-TIME-N, STOP-TIME-N)

Accepts up to five start/stop pairs:
 START-TIME-i, STOP-TIME-i
 where each component of the pair is real-14, to further delimit the action of debug statements. The limiting action becomes effective as soon as the program determines "what time it is". This command must be accompanied by some form of the DEBUG command or no debug output will be produced. The time frames are independent of, though not unaffected by, any other user selected timeframes.
- o TITLE("NEW-TITLE")

Accepts the string "NEW-TITLE", string-63, as the new banner to head all reports.
- o TRACES(SENSE-OF-RESTRICTION, TRACE-1,...,TRACE-N)

Excludes from, or includes in, the formatted dump only those trace subtypes listed. Exclusion is obtained when the SENSE-OF-RESTRICTION is DONTDO, alphanumeric-6; inclusion, when the SENSE-OF-RESTRICTION is DOONLY, alphanumeric-6. As many traces as required may be specified by TRACE-i, integer-3, in one or several commands; however, previous commands are discarded if the SENSE-OF-RESTRICTION changes from DONTDO.

15.6.3 Examples. The following are samples of the user optional commands:

NEW(1234)	Begins a fresh run with tape #1234 as lead tape.
DEBUG	Turns all debug statements on.
DEBUG(TIMSET)	Turns on debug statements in subroutine timeset.
HISTOG(1 ON)	Turns trailer flag on for histog 1, also turns on 1.
HISTOG(2, , 43, 1, 2)	43 buckets, lowval is 1, interval size is 2, report is on.
CPRINT(123)	123 seconds between periodic reports & plot lines.
LINES(DOONLY,"A4","T5","3225")	Format dump only lines with IDs a4 , t5 and the DRUN with ID 3225. Multiple "lines" are cumulative. Command also applies to debug during the analysis phase.
ON(1, 2 3)	Turns on reports 1, 2 and 3. (All reports are on by default).
PLOT(10, 200)	Plot 10 will stop after 200 points.
OPTION(-SUBDISPATCH)	Subtract users subdispatch time from his computed response time.
OPTION(MINUSDISKIO)	Subtract users disk I/O time from his computed response time.
OPTION(ONCHANGE)	Limits the TSS Delay - User Delay Report.
STOP(1234567)	Stop after processing 1234567 physical records.
TRACES(DONTDO, 1 9 78 43 66)	Do not format traces 1 9 78 43 and 66.
TRACES(DOONLY, 12)	Format only trace 12. "Dontdo" may be preceeded by a "doonly". Multiple "traces" are cumulative. Command also applies to debug during the analysis phase.
TIMEFR(0, 0200.)	Data reduction starts at 0200:00 with no stop time.
TIMEFRAME(10 234.58 0345.40)	Plot is bounded by 0234:58 and 0345:40.
TIMEFR(0, 0234.58, 0345.4)	Same for data reduction.
TIMEFR(DEBUG, 1142., 1150.)	If debug is on, then debug statements are bounded by 1142:00 thru 1150:00 independent of main timeframes.
ON OFF(1)	Two commands: turn on all reports except 1.
NEW(5677)	Must be last command on card.
TITLE("TSS reduction, version 7.2 - 1.000, 19 July 82")	Default title.

15.6.4 Card Input Errors. If the card input interpreter cannot understand a command keyword or data item, it will echo the entire card and mark the last correctly interpreted character position with the message "ILLEGAL COMMAND (or DATA) FOLLOWING...". Depending on the state of the interpreter, it may attempt to find the end of the incorrect command in order to pass on

to a following command on the same card, or it may reject the remainder of the card; it will report which of the two actions was taken, and will then continue with card interpretation. In no case will an action confirmed by the interpreter be reversed or "undone" by a following error.

15.7 JCL for Reduction Execution.

The data reduction procedures consist of a single FORTRAN program having a main level and multiple links.

A sample JCL deck is presented in figure 15-30. The LIMITS card should be studied to meet user needs. The run time (999) and output limit (60K) may both be altered as required by the duration of the monitoring run. As a rule of thumb, run time will be approximately one-to-one with data collection elapsed time if all three passes are made in the same activity.

TEARS requires 44K of memory in order to execute plus an additional 2K for SSA space.

If no data cards are required, the JCL card

```
$    file    05,null
```

should be used in their stead.

15.8 Multireel Processing

If the GMC collected data is on more than one tape reel, TEARS will coordinate the change with a series of messages to the computer operator. These messages first warn the operator of an impending change, then direct the change, and finally, tell the operator of the tape number error if it has occurred. You can help smooth the transition from reel to reel by including MSG2 cards in the JCL deck listing the tapes which comprise the data set. Here are the messages directed to the computer operator:

- a. PLEASE GET TAPE # yyyyy. IT WILL BE CALLED FOR SOON.

This warns of the impending change, and names the next tape in sequence (yyyyy).

- b. DISMOUNT OLD TAPE; --THEN MOUNT TAPE # yyyyy ON DRIVE zzzzz.

This directs the change, naming the new tape number (yyyyy) and the specific tape drive (zzzzz).

- c. WRONG TAPE JUST MOUNTED; CAN YOU MOUNT yyyyy ON DRIVE zzzzz (Y/N).

This message is conditional -- it is sent when TEARS finds the wrong tape has been mounted (by comparing internally generated tape

```

##a,j
$ ident 1820251/30/1878,prc-krsch
$ msg2 *****
$ msg2 ** please give it a "u" on tape error.
$ msg2 ** will use tapes jal76 and jbl76 twice
$ msg2 *****
$ select b29idpx0/object/tears
$ limits 999,44k,,60k
$ tape 01,x3dd,,jal76,,gmf-t90,,denl6
option(-subdispatch) option(-diskio)
cprint(20)
new(jal76)
on(emulate)
option(onchange)
$ endjob

```

Figure 15-30. TEARS JCL

labels), or when more subtle changeover errors cause TEARS to suspect the new tape is not correct. If the operator answers affirmatively, the message sequence begins again with message (b) above. If the answer is negative, TEARS will terminate the reduction pass gracefully. Otherwise, TEARS will repeat the message.

15.9 Tape Error Aborts

During the course of processing, the operator may be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up section and produce all reports generated prior to the tape error.

DISTRIBUTION

Addressee

CCTC Codes

C124 (Reference and Record Set)	3
C430	240
DCA Code 308c	1

Defense Technical Information Center, Cameron Station, Alexandria, Virginia 22314	<u>12</u>
	256

THIS PAGE INTENTIONALLY LEFT BLANK

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER OSM UM 246-82 - CH-3	2. GOVT ACCESSION NO. 10-4124 071	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Generalized Monitoring Facility Users Manual 1 May 1982		5. TYPE OF REPORT & PERIOD COVERED
7. AUTHOR(s) Barry Wallack George Gero		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer Performance Evaluation Branch Defense Communications Agency/COTC/CT51 Pentagon, Washington, D.C. 20301		8. CONTRACT OR GRANT NUMBER(s)
11. CONTROLLING OFFICE NAME AND ADDRESS Same as above		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Computer Performance Evaluation
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE 1 May 1982
		13. NUMBER OF PAGES 176
		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public Release -- Distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Honeywell, Performance Evaluation, Software Monitor, Resource Monitor, Tuning		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This Users Manual provides the information necessary for an individual to run all of the programs that are comprised in the Generalized Monitoring Facility. Program interrelationships are presented, as well as a general overview of the processing, input, and output procedures for each program. The Generalized Monitoring Facility is a sophisticated software monitor for the Honeywell 6000 computer systems.		

DD FORM 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

17-1

CH-3

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)



17-2

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

CH-3

FILM
2-8